

Everything You Always Wanted to Know About Optical Networking – But Were Afraid to Ask

Richard A Steenbergen <ras@petabitscale.com>

Updated: October 9, 2023

Purpose of This Tutorial

Why give a talk about optical at a routing/packet conference?

- Everything we do on “The Internet” largely revolves around fiber.
- Yet many router jockeys don’t get enough exposure to it.
- This leads to a wide variety of confusion, misconceptions, and errors when working with fiber optic networks.

Will this presentation make me an optical engineer?

- Probably not, and remember, I omitted almost all the math.
- The purpose of this tutorial is to touch on a bit of every topic, from the mundane to the advanced and unusual.
- But it helps to have a basic understanding of how and why things work, even if you aren’t designing fiber networks yourself.

The Basics of Fiber Optic Transmission

What is Fiber, and Why Do We Use It?


Fiber is ultimately just a “waveguide for light”.

- Basically: light that goes in one end, comes out the other end.
- Most commonly made of glass/silica, but it can also be plastic.

So why do we use fiber in the first place?

- Glass is an extremely low-cost material to produce and install.
- Very light (relative to copper), flexible, high-density material.
- Carries tremendous amounts of information (76.8 Tbps+ today).
- It can easily carry large numbers of completely independent signals over the same fiber strand, without interference.
- Can carry signals thousands of kilometers without regeneration.
- Technology continues to radically improve what we can do with our existing fiber infrastructure, without digging or disruption.

Modern High-Density Fiber Cables



Hold It Down Like I'm Giving Lessons in Physics

Light propagating through a vacuum is (theoretically) the maximum speed at which anything in the universe can travel.

- That speed is 299,792,458 meters per second, written as “ c ”.
- For shorthand math, round this up to 300,000 km/s.


But when light passes through materials that **aren't** a perfect vacuum, it propagates much slower than this.

- The speed of light in any non-vacuum material is expressed as a ratio relative to “ c ”, known as that material's “refractive index”.
- Water has a refractive index of “1.33”, or 1.33x slower than “ c ”.
- And when light tries to pass from one medium to another with a different index of refraction, a reflection can occur instead.
 - This is why you see a reflection looking up from under water.

Fiber Works by “Total Internal Reflection”


Fiber optic cables are internally composed of two layers.

- A “core”, surrounded by a different material known as the “cladding”.
- The cladding always has a higher “index of refraction” than the core.



- When the light tries to pass from the core to the cladding, and the angle is correct, it is reflected back into the core.

Demonstration Using a Laser Pointer



Duplex vs Simplex Operation

- The vast majority of deployed fiber optic systems are “duplex”.
 - One “pair” of fiber is used for each point-to-point link.
 - One strand is used to transmit a signal, the other to receive one.
 - This results in the simplest/cheapest optical components.
- But fiber itself is perfectly capable of carrying many signals, in both directions, over a single strand.
 - It just requires more expensive optical components to do so.
- Which system we use is largely a question of economics.
 - When fiber is cheap, we typically just throw burn extra strands.
 - But where the fiber is expensive or constrained (e.g. out of riser capacity, long-haul, etc) we use more advanced techniques.

The Most Basic Distinction in Fiber: Multi-Mode vs Single Mode

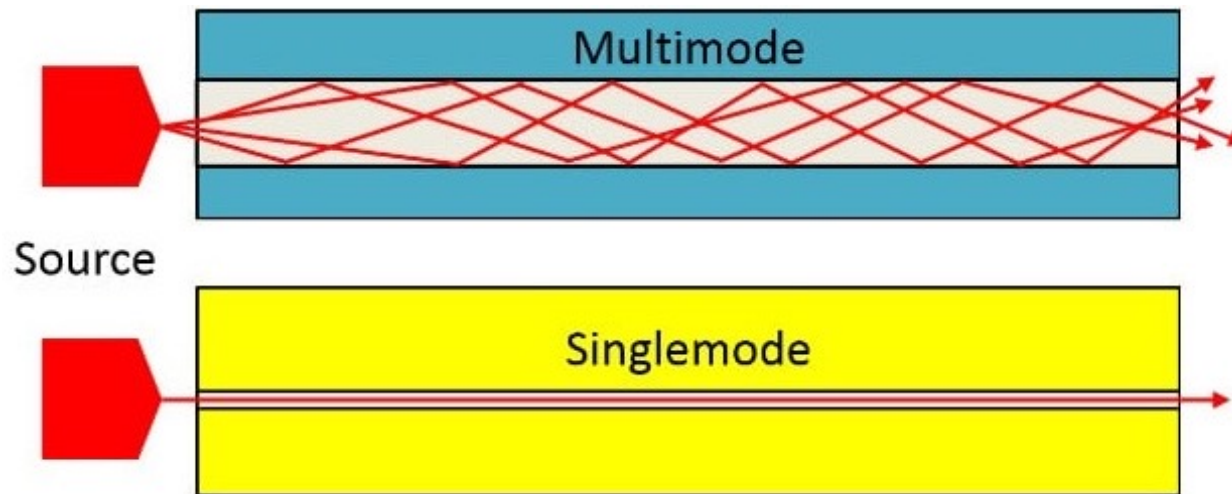
Multi-Mode Fiber (MMF)

- Specifically designed for use with “cheaper” light sources.
 - Has an extremely wide core, allowing the use of less precisely focused, aimed, and calibrated light sources.
 - But this comes at the expense of long-distance reach.
 - Fiber is so named because it allows multiple “modes” of light to propagate.
 - “Modal dispersion” typically limit distances to “tens to hundreds” of meters.
- Types of Multi-Mode
 - OM1/OM2 aka “FDDI grade”: found with orange fiber jackets.
 - OM1 has a 62.5 micron (μm) core, OM2 has a 50 μm core.
 - Originally designed for 100M/1310nm signals, starts to fail at 10G speeds.
 - OM3/OM4 aka “laser optimized”: found with “aqua” fiber jackets.
 - Specifically designed for modern 850nm short reach laser sources.
 - Supports 10G signals at much longer distances (300-550m, vs 26m on OM1).
 - Required for 40G/100G signals (which are really 10G/25G signals), 100-150m.


Single Mode Fiber (SMF)

- The fiber used for high bandwidths, and long distances.
 - Has a much smaller core size, between 8-10 microns (μm).
 - No inherent distance limitations caused by modal dispersion
 - Typically supports distances of $\sim 80\text{km}$ without amplification.
 - With amplification, can transmit a signal several thousand kilometers.
- SMF has an even broader array of types than MMF.
 - Also has “OS1” and “OS2”, but they’re packaging, not fiber type.
 - OS1 “tight buffered” for indoor use, OS2 “loose” to be blown into ducts.
 - “Classic” SMF can be called “SMF-28” (a Corning product name)
 - But it comes in many different formulations of Low Water Peak Fiber (LWPF), Dispersion Shifted Fiber (DSF), Non-Zero Dispersion Shifted Fiber (NZDSF), Bend Insensitive Fiber, etc.

Single Mode Fiber vs Multi-Mode Fiber



Optical Fiber Core Diameters



Fiber Cables and Connectors



Basic Optical Terms and Concepts

Optical Power and Attenuation

Quite simply, the brightness (or “intensity”) of the light.

- As light travels through fiber, some of its energy is lost.
 - Scattered by microscopic imperfections in the fiber.
 - Or absorbed by residual OH⁺ or dopants and converted to heat.
- This loss of intensity is called “attenuation”.
 - No matter how good the fiber, eventually the light will be lost.
- We typically measure optical power in “Decibels”
 - A decibel (dB, 1/10th of a Bel) is a logarithmic-scale unit expressing the relationship between two values.
 - The decibel is a “dimensionless-unit”, meaning it does not express an actual physical measurement on its own.

Optical Power and the Decibel

A decibel is a logarithmic ratio between two values

- -10dB is $1/10^{\text{th}}$ the signal, -20dB is $1/100^{\text{th}}$ the signal, etc.
- Another easy one: +3dB is double, -3dB is half.
- But remember, this doesn't tell you "double of what?"

To express an absolute value, we need a reference.

- In optical networking, this is known as a "dBm".
 - That is, a decibel relative to 1 milliwatt (mW) of power.
- 0 dBm is 1 mW, 3 dBm is 2 mW, -3 dBm is 0.5mW, etc.
- What does this make 0mW? Negative Infinity dBm.

Confusion between dB and dBm is one of largest sources of mistakes in optical networking!!!

Why Do We Measure Light in Decibels?

- Light, like sound, follows the “inverse square law”.
 - The signal is “inversely proportional to the distance squared”.
 - A signal travels distance X and loses half of its intensity.
 - The signal travels another distance X and loses another half.
 - At $2X$, only 25% remains.
 - At $3X$, only 12.5% remains.
 - At $4X$, only 6.25% remains.
- A logarithmic scale vastly simplifies our calculations.
 - A 3dB change is approx. half/double the original signal.
 - In the example above, there is a “3dB loss” per distance X .
 - At distance $2X$ there is 6dB of loss, at distance $3X$ it is 9dB, etc.
- This lets us to use elementary school math.
 - Which is much easier for the humans to work with.

Decibel to Power Conversion Table


Table 1 - Decibel to Power Conversion

dB (loss)	Power Out as a % of Power In	% of Power Lost	Remarks
1	79%	21%	---
2	63%	37%	---
3	50%	50%	1/2 the power
4	40%	60%	---
5	32%	68%	---
6	25%	75%	1/4 the power
7	20%	80%	1/5 the power
8	16%	84%	1/6 the power
9	12%	88%	1/8 the power
10	10%	90%	1/10 the power
11	8%	92%	1/12 the power
12	6.3%	93.7%	1/16 the power
13	5%	95%	1/20 the power
14	4%	96%	1/25 the power
15	3.2%	96.8%	1/30 the power
16	2.5%	97.5%	1/40 the power
17	2%	98%	1/50 the power
18	1.6%	98.4%	1/60 the power
19	1.3%	98.7%	1/80 the power
20	1%	99%	1/100 the power
25	0.3%	99.7%	1/300 the power
30	0.1%	99.9%	1/1000 the power
40	0.01%	99.99%	1/10,000 the power
50	0.001%	99.999%	1/100,000 the power

Fiber Optic Transmission Bands



- 850nm –
 - Highest attenuation, only used for short-reach applications.
- 1310nm – O-Band (“Original Band”)
 - Covers 1260nm-1360nm, 1310nm is the center frequency.
 - Point of lowest dispersion on classic SMF, but higher attenuation.
 - Primarily used for intermediate-reach applications (2-10km) today.
- 1550nm – C-band (“Conventional Band”)
 - Covers 1530nm – 1565nm, 1550nm is approx center frequency.
 - Lowest rate of attenuation, but higher chromatic dispersion.
 - Used for almost all long-reach and DWDM applications today.
- L-band (“Long Band”)
 - Covers 1565nm – 1625nm, 60nm of total spectrum.
 - Acts as an “extension” to C-band, with many similar properties.
 - Different amplification techniques are required.
 - Typically only used when fiber is extremely constrained/valuable.

Fiber Optic Transmission Bands



Impairments: Dispersion


- Dispersion simply means “to spread out”.
 - In optical networking, this results in signal degradation.
 - Eventually, it is no longer distinguishable as individual pulses at the receiver.
- We’re primarily concerned with two types of dispersion:
 - “Chromatic Dispersion” (CD), and “Polarization Mode Dispersion” (PMD).



Chromatic Dispersion (CD)

- Different frequencies of light propagate through a non-vacuum at different speeds.


- This is how rainbows and optical prisms work.
- Eventually, your signal stops arriving at the same time, and become indistinguishable.



- Chromatic Dispersion has historically been a huge limiting factor.
 - The faster your symbol rates, the wider your signal linewidths.
 - The wider the signal linewidth, the worse CD affects your signal.
 - CD impacts increase at the square of the baud rate.
 - We measure CD in “picoseconds per nanometer” (ps/nm).

Polarization Mode Dispersion (PMD)

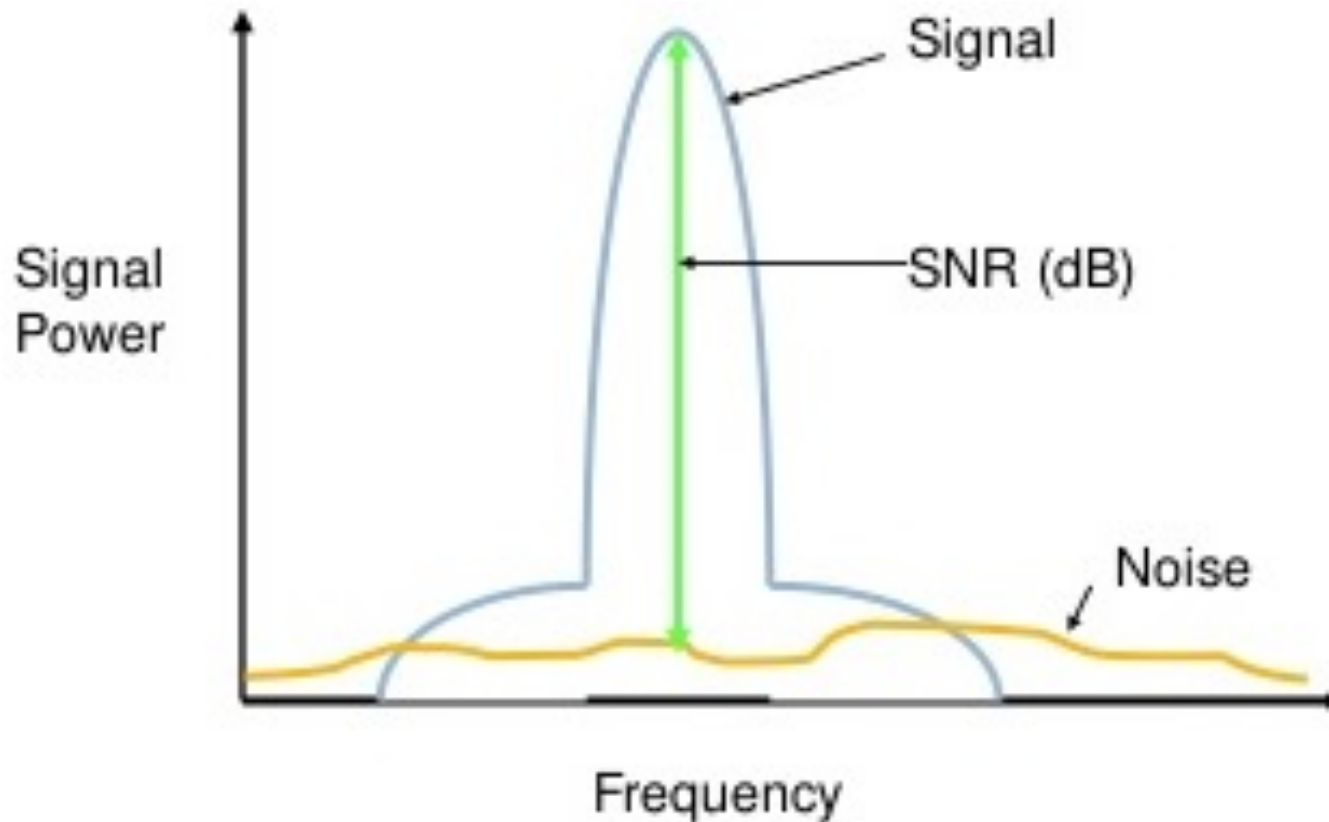
- Polarization Mode Dispersion (PMD) results when the light of one polarization arrives at a different time than the other.
 - This is caused by imperfections in the cylindrical shape of the fiber.



- The difference in arrival time between the polarizations is called “Differential Group Delay” (DGD), measured in picoseconds (ps).

Optical Signal to Noise Ratio (OSNR)


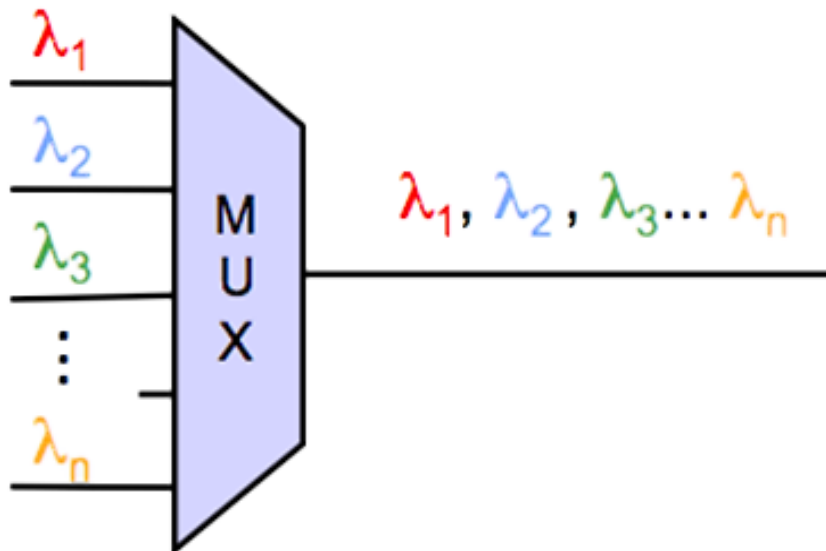
In amplified long-reach systems, OSNR becomes the single most important limiting factor.



Wave Division Multiplexing

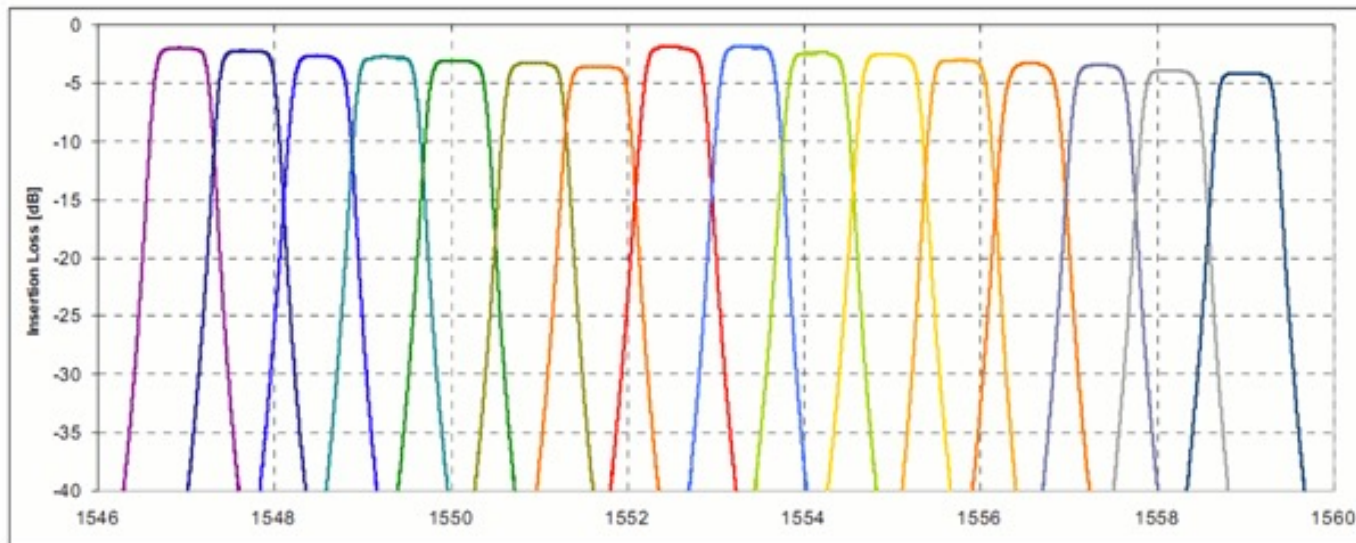
Wave Division Multiplexing (WDM)

- We know that light comes in many different “colors”.
 - What we perceive as “white” is just a mix of many wavelengths.
- These different colors can be combined on the same fiber.
- The goal is to put multiple signals on the same fiber without interference (“ships in the night”), thus increasing capacity.



Coarse Wave Division Multiplexing (CWDM)

- CWDM is loosely used to mean “anything not DWDM”
 - One “popular” meaning is 8 channels with 20nm spacing.
 - Centered on 1470 / 1490 / 1510 / 1530 / 1550 / 1570 / 1590 / 1610




- With Low Water Peak fiber, another 10 channels are possible
 - Centered on 1270/1290/1310/1330/1350/1370/1390/1410/1430/1450.
- Can also be used to refer to a simple 1310/1550nm mux.

Dense Wave Division Multiplexing (DWDM)

- Defined by the ITU-T G.694.1 as a “grid” of specific channels.
- Within C-band, the follow channel sizes are common:
 - 200GHz – 1.6nm spacing, 20-24 channels in C-band
 - 100GHz – 0.8nm spacing, 40-48 channels in C-band
 - 50GHz – 0.4nm spacing, 80-96 channels in C-band
 - 25GHz – 0.2nm spacing, 160-192 channels in C-band
- A rough guideline to the technology:
 - 200GHz is “2000-era” old tech, rarely seen in production today.
 - 100GHz is still quite common for metro DWDM tuned pluggables.
 - 50GHz is common for commercial, long-haul, and 100G systems.
 - 25GHz was used briefly for high-density 10G systems, before the move to coherent 100G systems shifted everything back to 50GHz.
 - Modern systems can be flexible, in 6.25GHz increments or smaller.

WDM Channel Size Comparison



What Are The Advantages to Each?

CWDM

- Cheaper, less precise lasers can be used.
 - The actual signal in a CWDM system isn't really any wider.
 - But the wide channel allows for large temperature variations.
 - Cheaper, uncooled lasers more easily stay within the window.


DWDM

- Far more channels are possible within the same fiber.
 - 96 channels (at 50GHz) in 32nm of spectrum, vs. 8ch in 160nm.
- Can stay completely within the C-band
 - Where attenuation and dispersion are far lower.
 - And where simple Erbium Doped Amplifiers (EDFAs) work.
- But can also be duplicated within the L-band for expansion.


Other Uses of Wave Division Multiplexing

But other forms of WDM exist as well

- Simple 1310/1550 muxes.
- 4-lane “Grey” Optics
 - New high-speed interfaces often start using WDM lanes.
 - They’re cheaper to implement, or can support older fiber technology.
 - 10GE had 10GBASE-LX4 (4x 2.5G channels, rather than 1x 10G)
 - 40GE has LR4 (4x 10G, 1270nm / 1290nm / 1310nm / 1330nm)
 - 100GE has LR4 (4x 25G, 1295nm / 1300nm / 1305nm / 1310nm)
- Single Strand Optics (BX “bidirectional” standards)
 - E.g. 1310 / 1490nm mux integrated into a pluggable transceiver.



DWDM Channel Terminology



WDM Networking Components

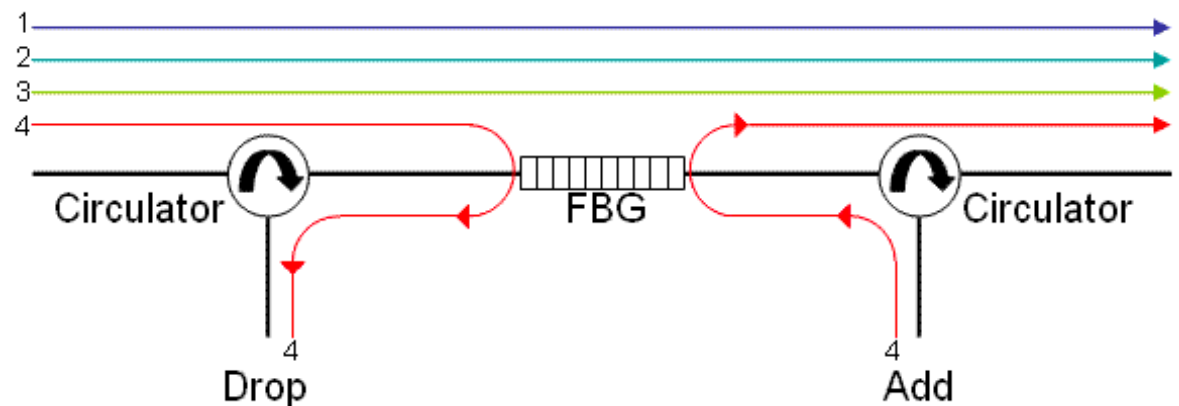
WDM Mux/Demux

- A simple, passive (unpowered) device, which combines/splits multiple colors of light to/from a single “common” fiber.
- Short for “multiplexer”, sometimes called a “filter”, or “prism”.
 - A “filter” is how it actually works, by filtering specific colors.
 - But people conceptually understand that a prism splits light into its various component frequencies.
- A complete system requires both a mux and a demux, for the TX and RX operation.
- Typically sold as a single package containing both units, for simplicity (and use on duplex fiber).







Optical Add/Drop Multiplexer (OADM)

- Selectively Adds and Drops certain WDM channels, while passing other channels through without disruption.
- While muxes often used at major end-points to break out all channels, OADMs are often used at mid-points within rings.
- A well-constructed OADM ring can reach any other node in the ring optically, potentially reusing some wavelengths multiple times on different portions of the ring.
- This is also an entirely passive and unpowered device.



Passive Optical Filter Technology

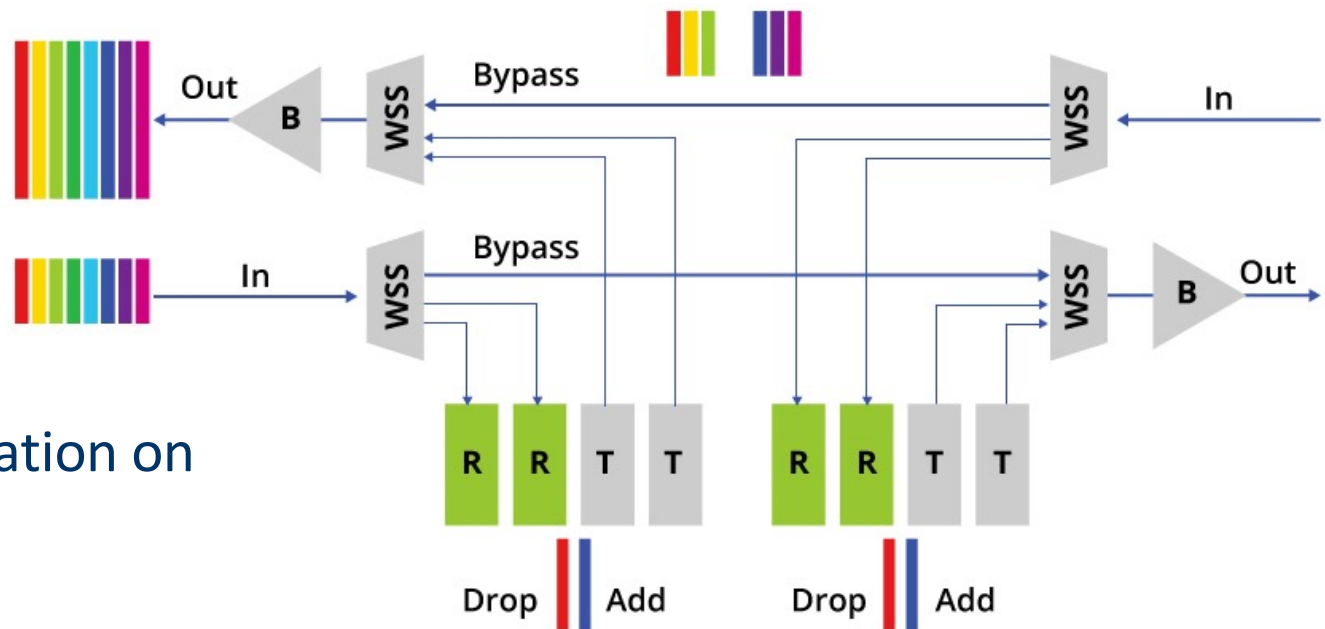
- FBG (Fiber Bragg Grating)
 - An “etched” fiber core, which causes a certain frequency to be reflected.
 - Does one channel at a time.
- Thin Film Filter (TFF)
 - Used for medium channel counts, < 16
- AWG (Arrayed Waveguide Grating)
 - Typically used for 40+ channel systems.
 - Essentially a very fancy interferometer.
 - For high channel counts, these are the cheapest and have lowest insertion loss.
 - Early versions required active cooling, but most units today are athermal (AAWG).



Reconfigurable OADM (ROADM)

A ROADM is a software reconfigurable OADM

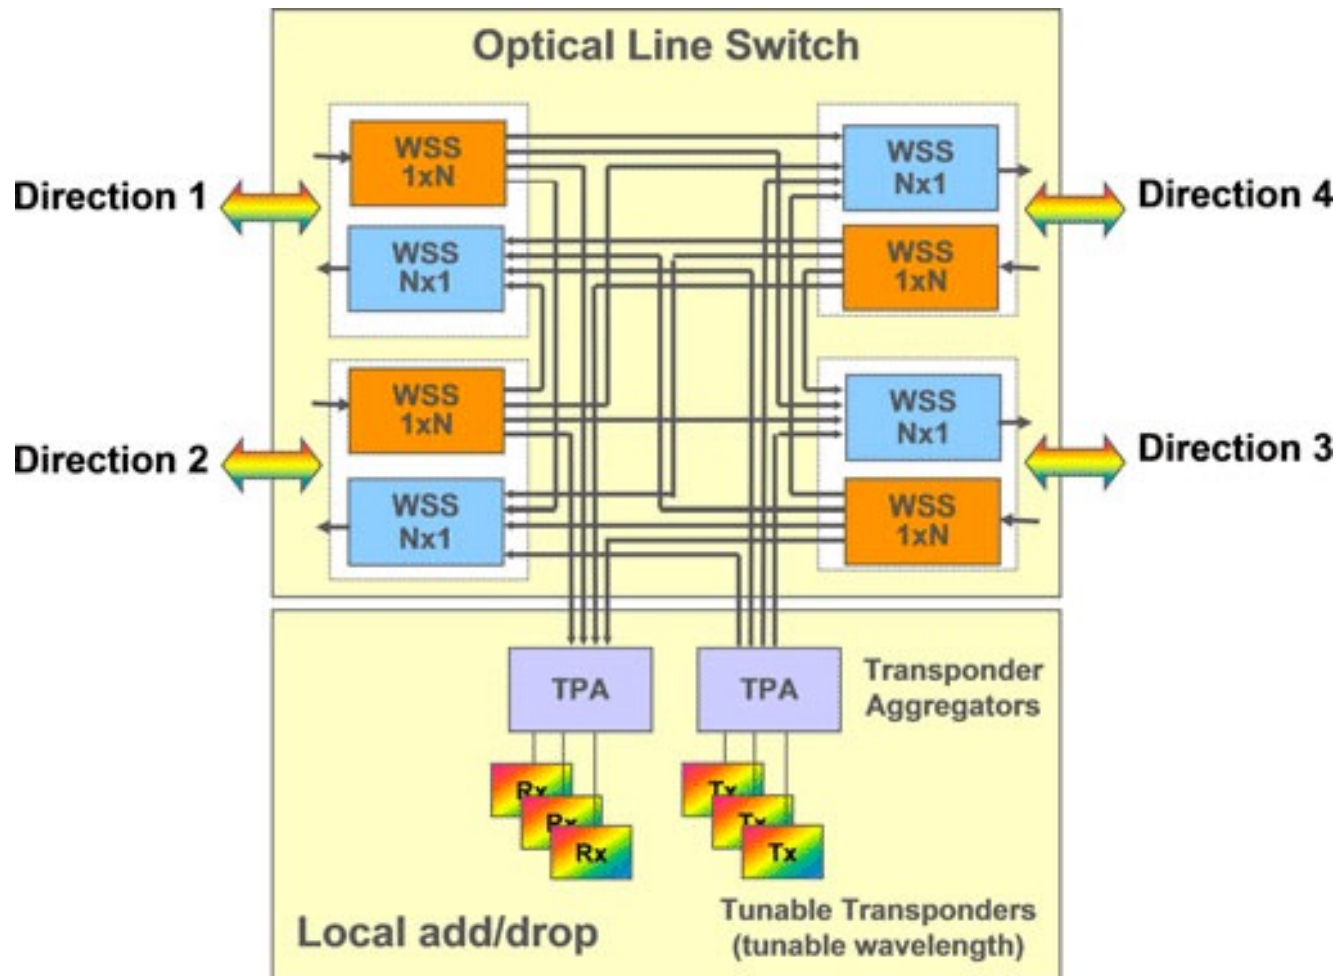
- Built around a WSS (Wavelength Selective Switch).
- Common Architectures:
 - Broadcast and Select: A 1xN Passive Splitter on TX, and WSS on RX
 - Route and Select: A WSS on TX, and WSS on RX



- Plus a lot of amplification on all sides!


ROADM Degrees

A ROADM “degree” is one possible direction light can be steered.




Fixed OADM (FOADM) vs ROADM Rings

Fixed OADM




- Optical connections between nodes are static, not reconfigurable without hw swap.
- Requires extensive man hours to support.
- But incredibly cheap and built on simple passive components.

ROADM




- Dynamically connects any point to any point optically, using only software.
- Can be run from a point-and-click GUI.
- But built on incredibly complex and expensive pieces of equipment.

The Evolution of the ROADM




Basic ROADM

- Reconfigurable, but add/drop still goes to a standard fixed mux.
- Specific frequencies must be connected to specific ports.
- The network must be recabled in order to change or move frequencies.



Colorless ROADM

- Eliminates the need to map specific frequencies to specific ports.
- But still limited to muxing in one direction at a time.



CDC ROADM

- Colorless – Any channel can be add/dropped on any port.
- Directionless – Any channel can be sent to any direction.
- Contentionless – The same channel can be reused in different directions without causing internal contention in the ROADM.

Modern Networking and the CDC ROADM

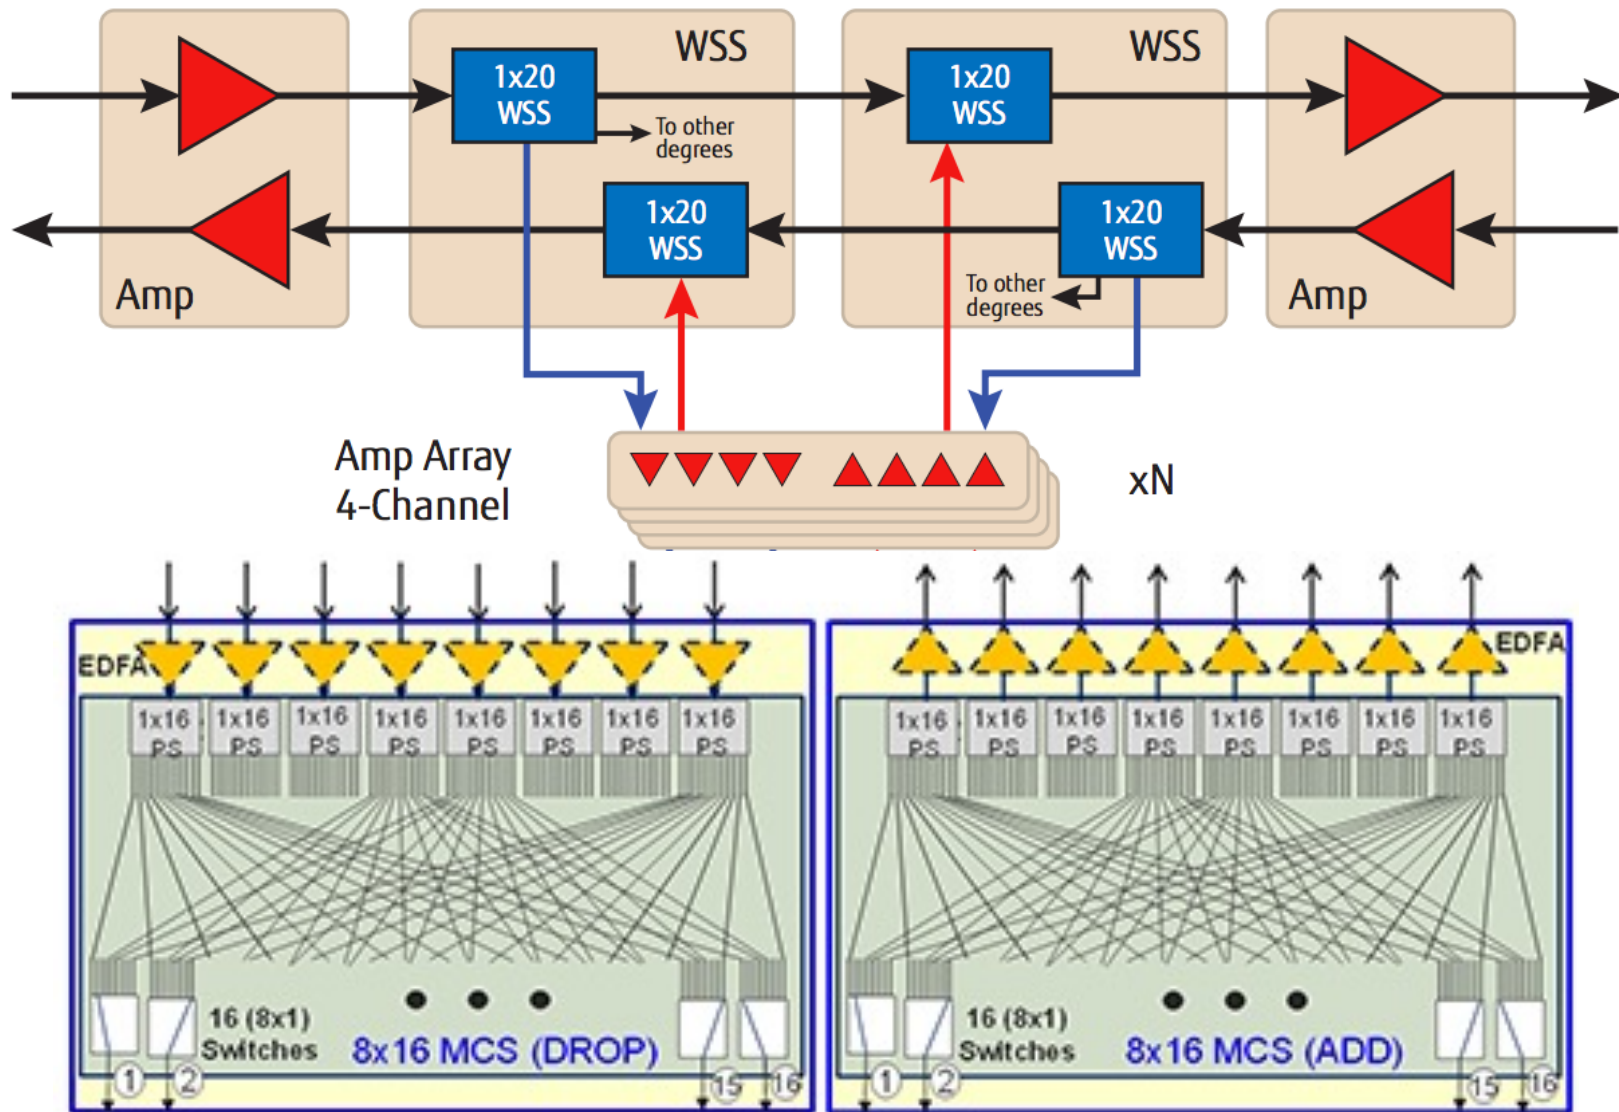
The goal is to move optical channels entirely in software.

- Transponders can be reallocated onto different physical paths as traffic patterns change (due to time-of-day changes, or during fiber cuts), potentially increasing efficiency and reducing costs.
- Eliminates the need to physically move cables to reconfigure.
- Allows dynamic bandwidth allocation at an optical level.

Potentially the entire process can be automated.

- IETF pushing for vendor interoperability and signaling via mechanisms like PCEP (Path Computation Element Protocol).
- Routers could “request” additional bandwidth from a pool of underlying transponders, based on real-time traffic requirements.


What Goes Into a Modern CDC ROADM





DWDM Superchannels

What if we want performance that a single carrier can't deliver?

- Superchannels pack multiple carriers together in a single channel.
- Often you can pack the carriers together tighter than if you were using standard channels too, increasing spectral efficiency.
- In this example, we deliver 1 Tbps using off the shelf 10x100G technology.



The Evolution of DWDM Channels



Other Optical Networking Components


Dispersion Compensation Units (DCU/DCM)

- Originally, just a big spool of fiber in a box.
 - Designed to cause dispersion in the opposite direction (opposite "slope") as the transmission fiber.
 - Passing through this spool reverses the effects of dispersion caused by the transmission fiber.
 - But it also adds fiber distance (typically 20%), and latency.
 - Usually deployed at amp sites.
- Newer designs are FBG based.
 - Reduces insertion loss, eliminates the latency overhead and physical bulk of a fiber spool-based design.
 - But also specific to channel plan, does not align well wider channels.





Optical Switches

- Optical Switches
 - Let you direct light between ports, without doing O-E-O conversion.
 - Built with an array of tiny mirrors, which can be moved electrically.
 - Allows you to connect two fibers together optically in software.
 - Becoming popular in optical cross-connect and fiber protection roles.



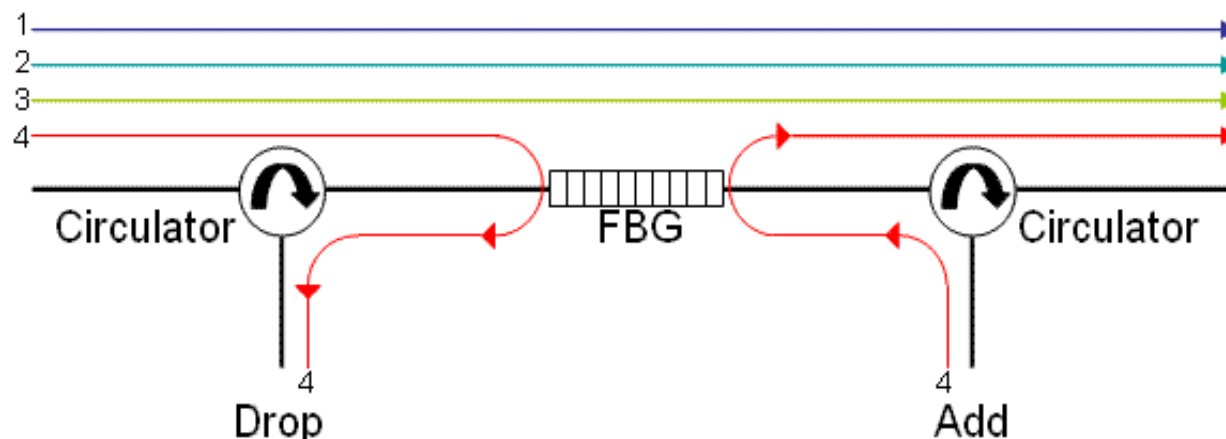
Wavelength Selective Switch (WSS)

- Lets you “route” an individual wavelength between ports
 - The WSS is a key component inside of a ROADM.
 - First generation WSS’ used 3D MEMS optical switches.
 - Modern WSS’ use Liquid Crystal on Silicon (LCoS).



Circulator

- A component typically not seen by the end user.
 - A circulator has 3 fiber ports.
 - Light coming in port 1 goes out port 2.
 - Light coming in port 2 goes out port 3.
 - Frequently used inside other popular components.
 - Bragg grating based components, like OADM's and small muxes.
 - Dispersion compensation spools, amplifiers, etc.
 - Very useful when building single-strand bidirectional systems too.



Optical Splitters


- Do exactly what they sound like they do, split a signal.
- Some common examples of splitters are:
 - A 50/50 Splitter
 - Often used for simple “optical protection”.
 - Split your signal in half and send down two different fiber paths.
 - Use an optical switch with power monitoring capabilities on the receiver, have it automatically pick from the strongest signal.
 - If the signal on one fiber drops, it switches to the other fiber.
 - A 99/1, 98/2, or 95/5 Splitter
 - Often used for “Optical Performance Monitoring”.
 - Tap a small % of the signal and send it to a spectrum analyzer.
 - A 1xN Splitter (e.g. 1x32)
 - Used in multipoint signal distribution, e.g. Cable TV

Optical Amplification

Optical Amplifiers


Optical amplifiers increase the intensity of a signal

- Purely optical way to extend signal reach, no regeneration.
- There are different types, for different frequencies of light.
- And different designs, for different positions within the span.
 - Booster Amplifiers are designed for high total output powers.
 - Pre-Amplifiers are designed for low input powers with minimal noise.
 - Inline Amplifiers strike a balance between the two.



Erbium Doped Fiber Amplifier (EDFA)

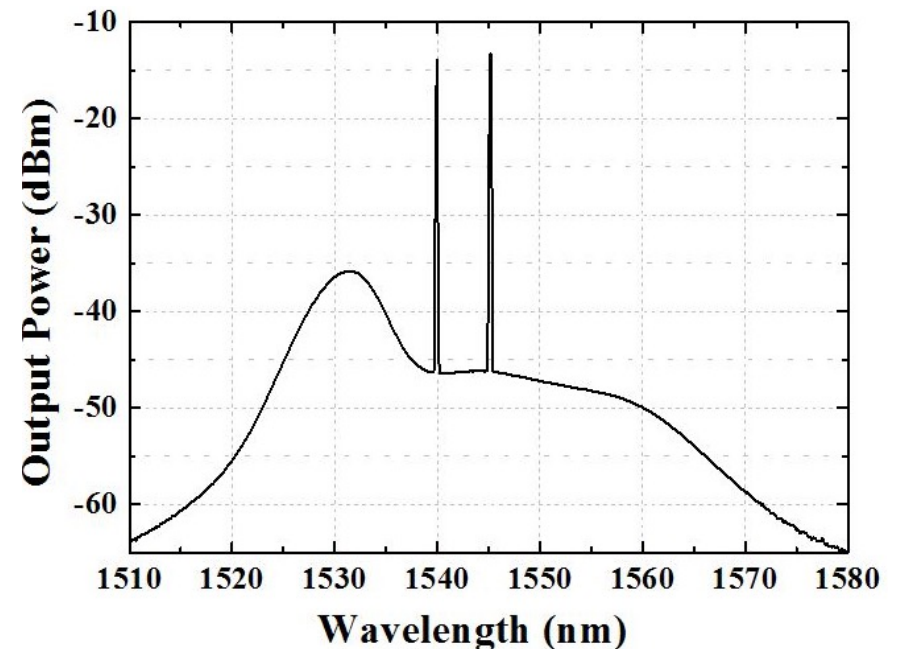
- The most basic/common fiber amplification system.
 - In an EDFA, a piece of fiber is “doped” with Erbium ions (Er^{3+}).
 - Another laser (980/1480nm) is pumped in via a coupler.
 - Pump laser puts Erbium electrons into higher energy state.
 - C-band frequency photons cause the Erbium electrons to decay to their ground state, and emits a clone of the original photon.



EDFA Noise

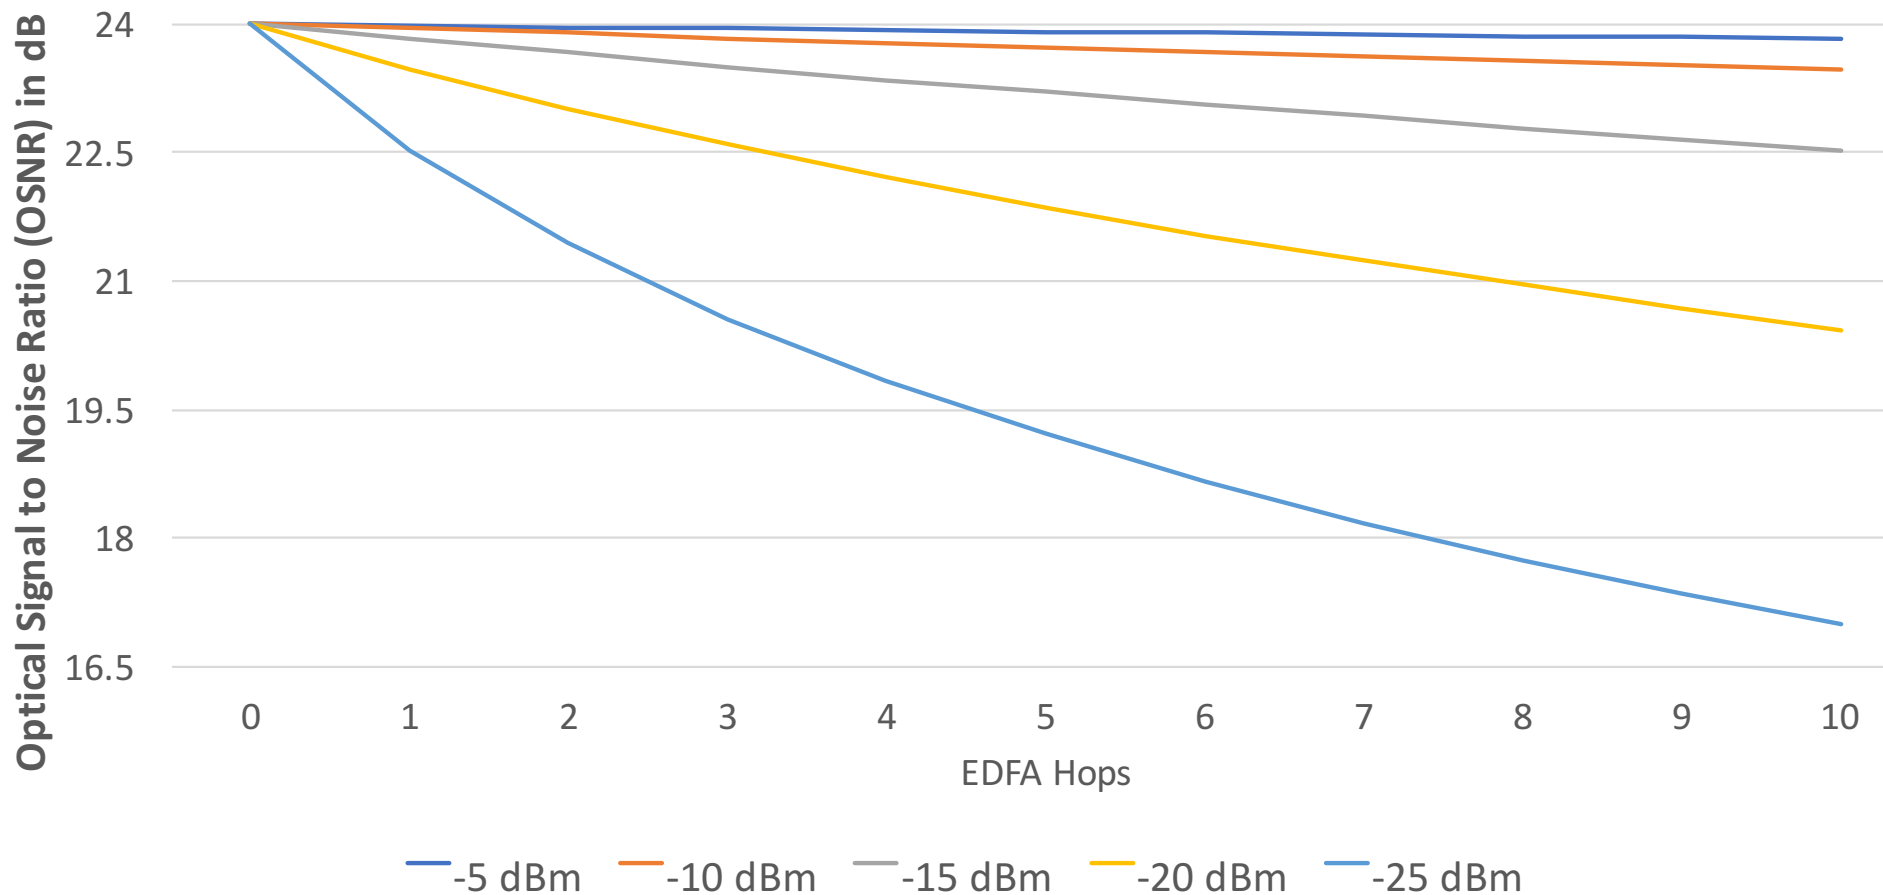
So why can't we use this to boost a signal forever?

- In addition to the intended boosting of signal, EDFAs also generate noise (“Amplified Spontaneous Emission”, ASE).
- Whenever an excited Erbium electron fails to encounter a “good” photon within $\sim 10\text{ms}$, it falls back to its ground state spontaneously, and emits a “noise” photon.
- Once generated, this noise is indistinguishable from the original signal.
- After enough amplifications, the noise ruins the original signal.



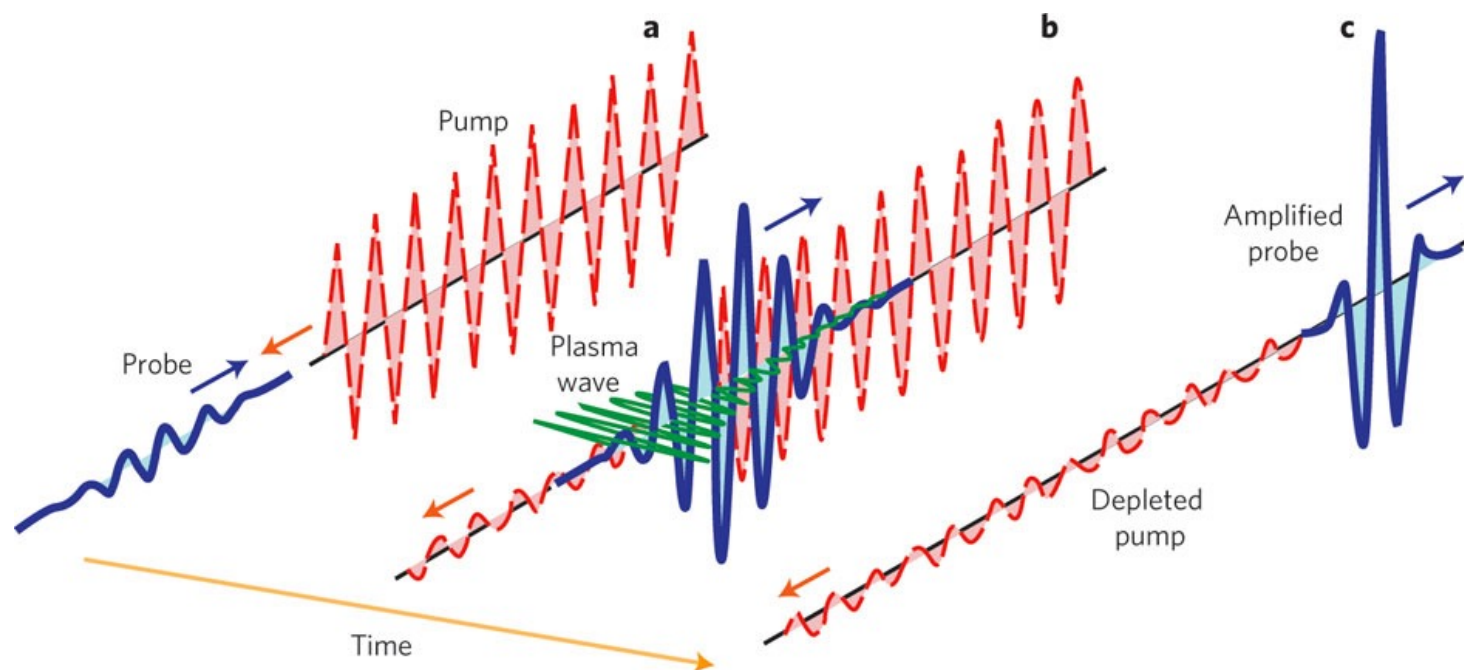
EDFA Noise – Why Input Power Matters

OSNR(dB) Following Cascaded EDFAs, at Various Input Power Levels

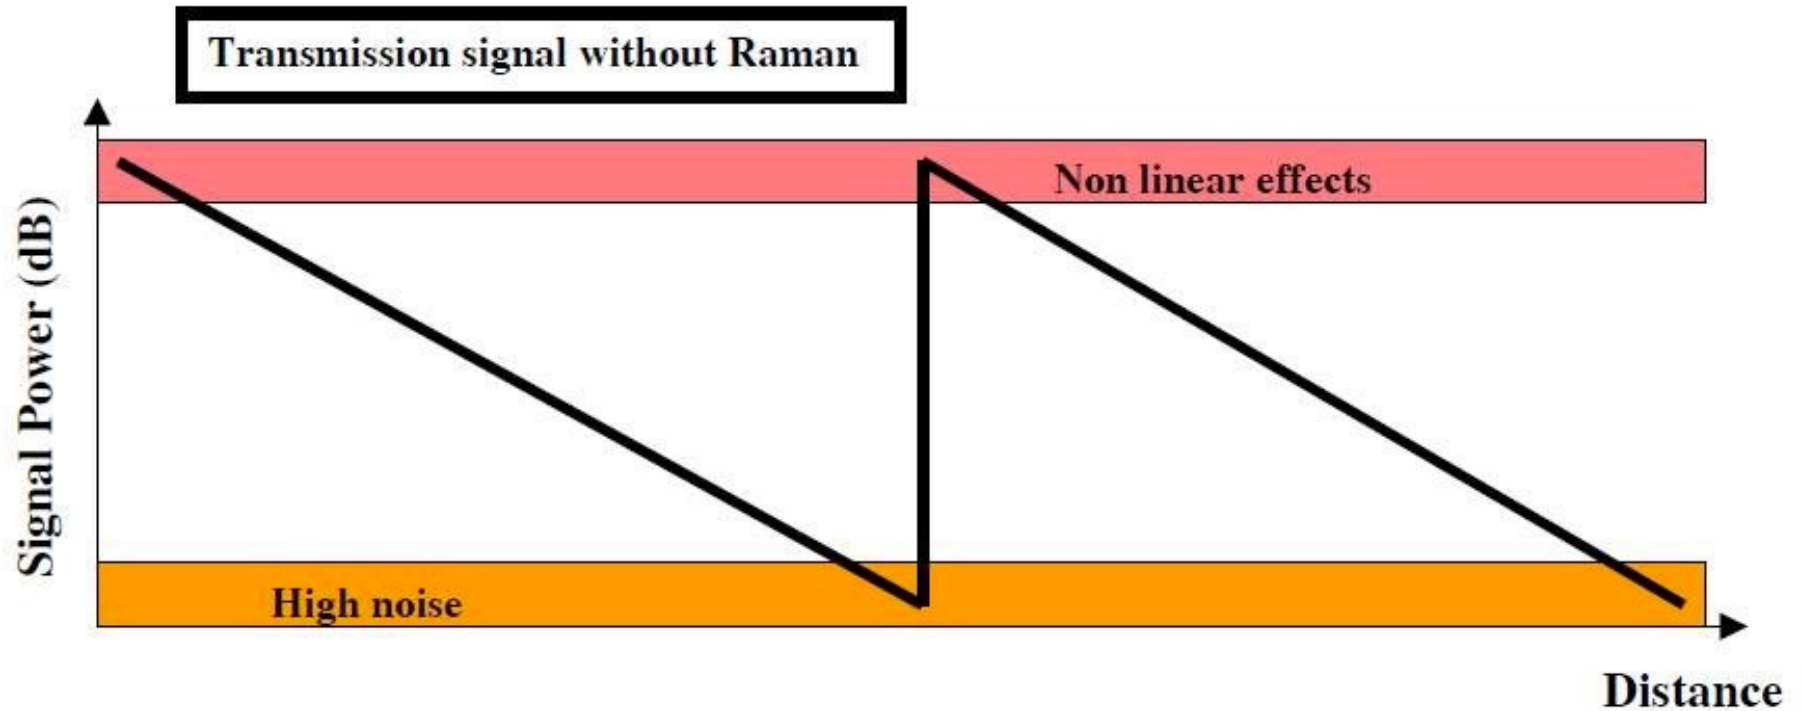



Raman Amplification


- Works on a principal of “Stimulated Raman Scattering”.
- Requires very high-power pump lasers, long gain mediums.
 - EDFAs used “lumped” design, gain media of 1-20 meters.
 - Raman usually use “distributed” design, gain medium 20+ km.
- Best used “counter-propagating” (facing backwards towards signal)



EDFA Only Amplification




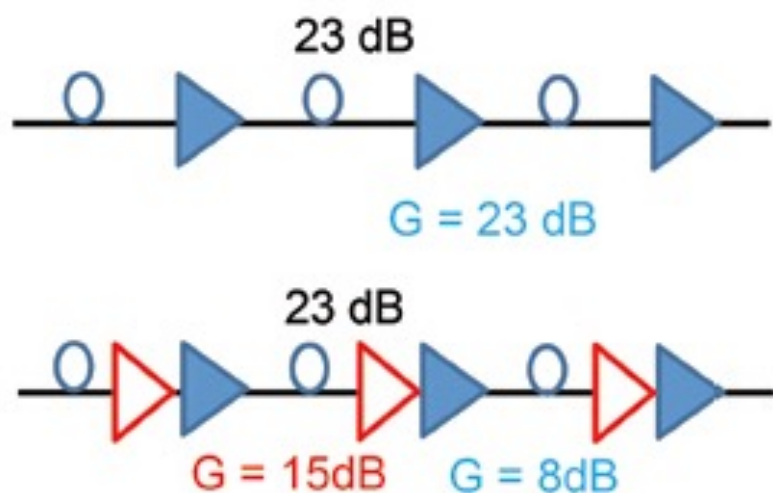
Hybrid EDFA + Raman Amplification



Hybrid EDFA + Raman Performance


Adding Raman extends EDFA reach significantly!

- In this example (21dB OSNR): from 7 hops to 20 hops
- At 100km/each, we go from only 700km to doing 2000km.



Amplifiers and Power Balance

- Amplifiers introduce some of their own unique issues
 - Unbalanced channel power causes OSNR penalties.
 - Amplifier gain often varies significantly across frequencies
 - Gain Flattening Filters try to compensate for this property.
 - Typical gain variations between channels (“ripple”) are still $< 1\text{dB}$.



- Even small power variations can add up after several hops.
 - Dynamic Gain Equalization (“DGE”) is required periodically.
 - ROADMs are often used in this role, to balance every channel.

Amplifiers and Total System Power


- Amplifiers also have limits on their total system power
 - Both what they can output, and what they can take as input.
 - And the total input power changes as you add channels!
 - A single channel at +0 dBm is only 1mW of input power.
 - 40 DWDM channels at +0 dBm/each is 40mW, or +16dBm of power.
 - If your amplifier's maximum input power is -6dBm, and you run 40 DWDM channels through it, each channel must be below -22dBm.
 - Failing to plan for this can cause problems as you add channels.
 - The total input power also changes as you lose channels.
 - Imagine power fails to a POP, and many channels are knocked offline.
 - Suddenly the total system power has changed significantly.
 - A good EDFA needs to constantly monitor and adjust power levels.
 - The best EDFAs will communicate with others on the line system.

Other Optical Networking Concepts

Forward Error Correction

- FEC adds extra/redundant information to the transmission, so the receiver can computationally “recover” from any errors.
- In practice, FEC works by lowering the required OSNR, which can help an otherwise unusable signal function normally.
 - Using clever math, padding a 10.325Gbps signal to 11Gbps (7% overhead) can extend a signal from 80km to 120km or beyond.
 - This can really matter when upgrading older DWDM systems.
 - Since it usually isn't practical to move amp sites closer on a live system.
 - FEC has evolved significantly as well.
 - 1st Gen – RS-FEC – 6% overhead for ~6dB of net coding gain.
 - 2nd Gen - EFEC – 7% overhead for ~8-9dB of net coding gain.
 - 3rd Gen – SD-FEC – 20-25% overhead for 10-11dB coding gain.
 - This might not seem worth it, but a 1-2dB gain in OSNR can add hundreds of kilometers of reach.
 - FEC is now an essential part of modern client optic standards as well.


The Benefits of Forward Error Correction



OTN Digital Wrapper Technology (G.709)

OTN stands for Optical Transport Network

- Replacement for SONET/SDH, with support for optical networking.
 - A standard for the generic transport of any protocol across a common optical network, with TDM mux/demux capabilities.
 - Implemented as a “wrapper” around other protocols.
- Why is this needed?
 - Pure optical channels only make sense for high-speed protocols.
 - Example: A single 100GE service, delivered over a 100G wave.
 - Low speed services still need to be aggregated.
 - Example: 10x10GE services on a 100G wave.
 - OTN lets the optical network be completely transparent to underlying protocols.
 - Can also help with troubleshooting.



Types of Single Mode Optical Fiber

Types of Single-Mode Fiber

- We've already discussed how single-mode fiber is used for essentially all long-reach fiber applications.
- But there are also many different types of SMF.
- The most common types are:
 - “Standard” SMF (ITU-T G.652) A.K.A. SMF-28
 - Full Spectrum (Low Water Peak) Fiber (ITU-T G.652.C/D)
 - Dispersion Shifted Fiber (ITU-T G.653)
 - Cutoff Shifted Low-Loss Fiber (ITU-T G.654)
 - Non-Zero Dispersion Shifted Fiber (ITU-T G.655)
 - Bend Insensitive Fiber (ITU-T G.657)


“Standard” Single-Mode Fiber (G.652)

- One of the original fiber cables.
 - Deployed widely throughout the 1990s.
- Frequently called “SMF-28”, or simply “classic” SMF.
 - SMF-28 is actually a Corning product name.
 - Also called NDSF (Non-Dispersion Shifted Fiber).
- Optimized for use by the 1310nm band.
 - Has the lowest rate of dispersion here.
 - Originally deployed before the adoption of WDM.
- Ironically, has come full circle to again being the best choice for modern high-speed DWDM systems.


Low Water Peak Fiber (G.652.C/D)

- Modified G.652, designed to reduce water peak.
 - Water peak is a high rate of attenuation at certain frequencies due to OH- hydroxyl molecule within the glass.
 - This high attenuation makes certain bands “unusable”.

Absorption of Light by Hydrogen at Various Wavelengths



Attenuation of Standard vs. Low Water Peak Fiber



Dispersion Shifted Fiber (ITU-T G.653)

- An attempt to improve dispersion at 1550nm.
 - The rate at which chromatic dispersion occurs changes depending on the frequency of light.
 - The point of lowest dispersion in G.652 occur around 1310nm.
 - But this is not the point of lowest attenuation, which is around 1550nm.
 - DSF shifts the point of lowest dispersion to 1550nm too.
- But this turned out to cause huge problems.
 - It worked well for simple, single channel systems.
 - But running DWDM signals over DSF causes huge amounts of non-linear interactions at higher powers.
 - As a result, this fiber is rarely used today.

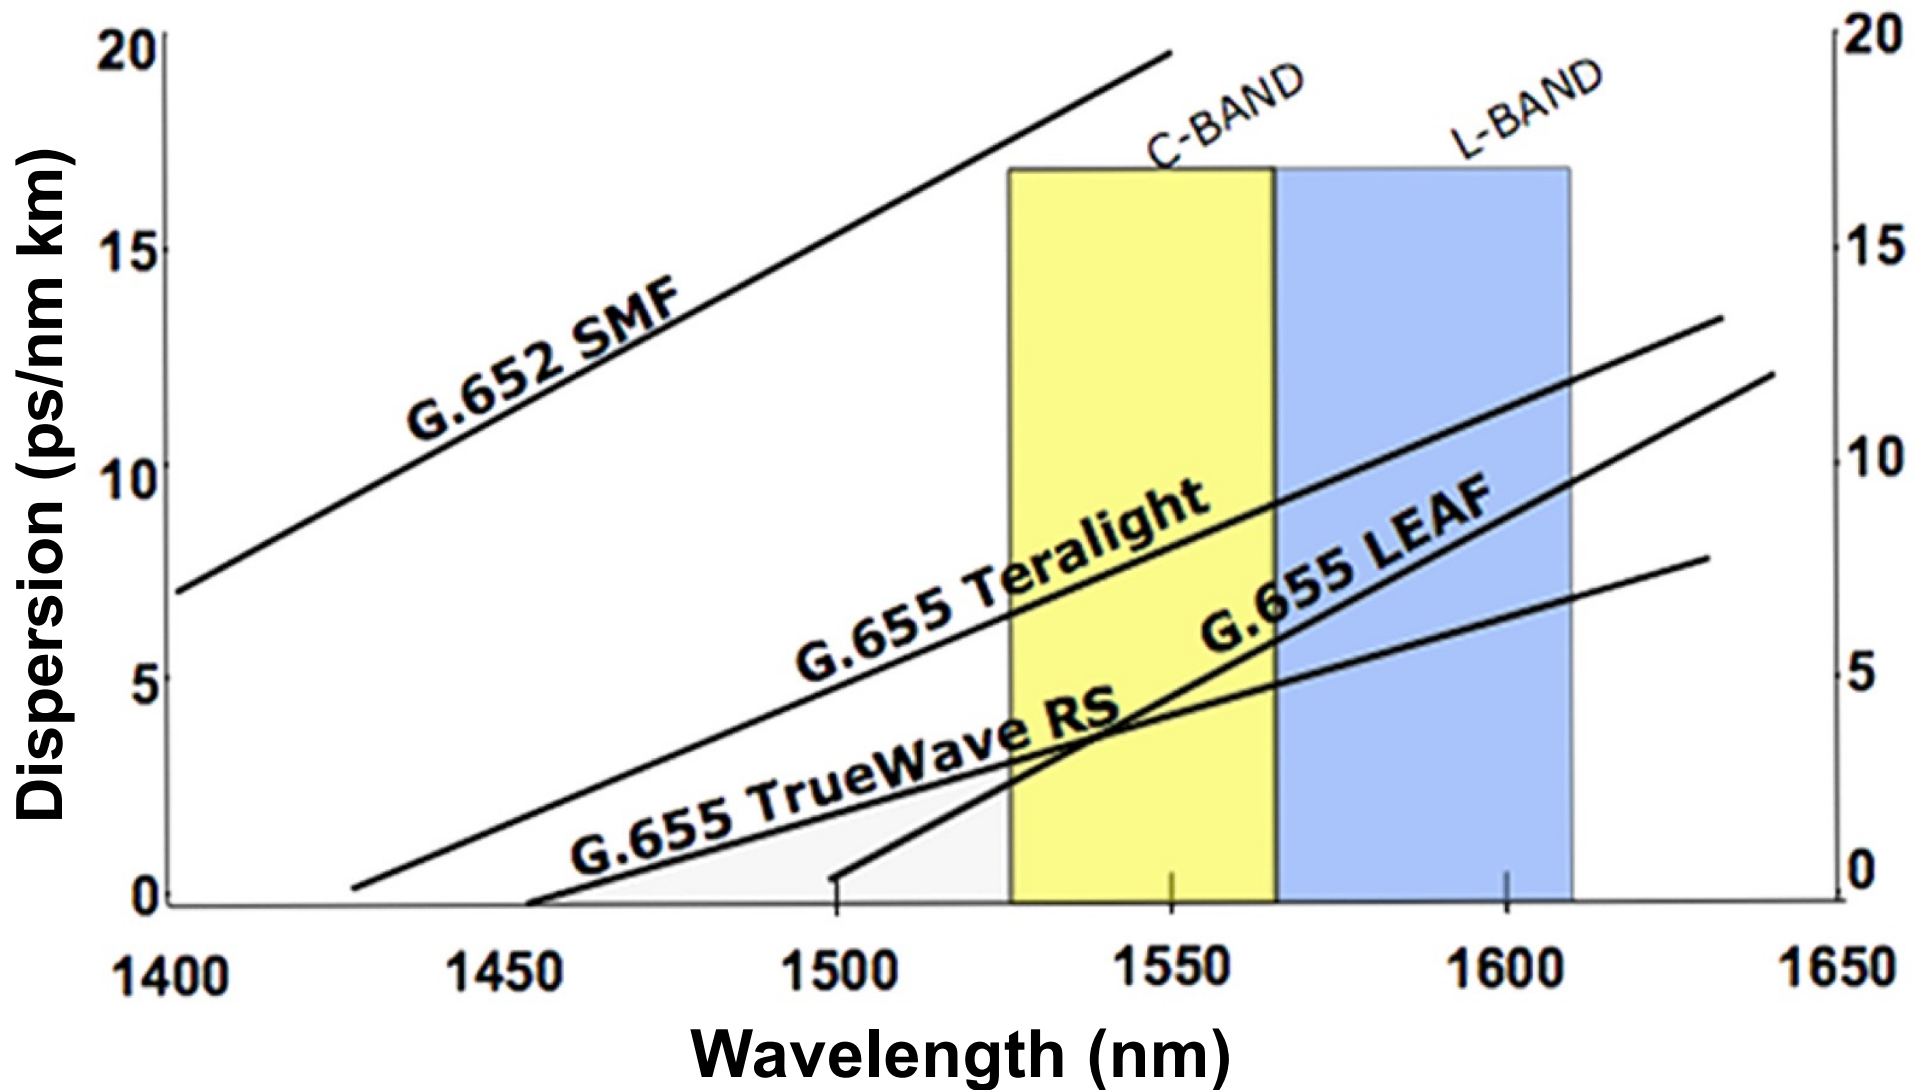
Non-Zero Dispersion Shifted Fiber (G.655)

- Similar concept to Dispersion Shifted Fiber
 - But the zero point is moved just outside of C-band.
 - This leaves a small amount of dispersion, which helps avoid the non-linear impairments caused by DSF.
- To manage dispersion, NZDSF comes in 2 types
 - NZD+ and NZD-, with opposite dispersion “slopes”.
 - The “transmission fiber” still spreads out 1550nm just a bit.
 - Then “compensation fiber” compresses it in the opposite direction.
 - By switching between the two slopes, the original signal can be maintained even over extremely long distances.

Other Single-Mode Fiber Types

- G.654
 - Ultra low attenuation, high power capable fiber.
 - Designed for ultra-long reach systems like undersea cables.
- G.657
 - Bend Insensitive fiber (reduced sensitivity at any rate).
 - Uses a higher refractive index cladding than normal fiber.
 - Designed for patch cable use, where a perfect bend radius may not be practical.
- Modern fibers are often better than these specs.
 - But much of what's actually in the ground is old fiber.

Dispersion Rates of Commercial Fibers



Non-Linear Impairments

Non-Linear Impairments

- Might be better described as “high power problems”.
 - If you don’t transmit at high powers, you’ll never see them.
 - But if you care about reach, you’ll probably be trying to push this.
 - What is “high power”? “Depends”, but usually above +4dBm / carrier.
- Non-Linear Impairments can be categorized as:
 - Stimulated Scattering
 - Stimulated Brillouin Scattering (SBS)
 - Stimulated Raman Scattering (SRS)
 - Kerr Effect
 - Intense light causes changes to the refractive index of the fiber.
 - Four Wave Mixing (FWM), Self-Phase Modulation (SPM), Cross-Phase Modulation (XPM)

Stimulated Brillouin Scattering (SBS)

- Excessive power transmitted into the fiber causes acoustic vibration at an atomic level within the lattice structure of the glass.
 - These vibrations set up Bragg grating effects, causing reflections.
 - Past a certain point, power is reflected back rather than forwards.
 - This limits power, causes errors, and can damage the transmitter.
- SBS is highly dependent on the “power density” in the fiber.
 - Wider linewidths spread the optical power out over more freq.
 - SBS suppression techniques include “dithering” to a wider signal.
 - Coherent helps quite a bit here, higher baud rates do too.
- SBS impact also largely requires long distances of fiber.
 - Putting high power through a very short span may not hurt you.
 - Typical “effective length” maxes out at around 20km.

Stimulated Raman Scattering (SRS)


- SRS is related to the SBS phenomenon.
 - Used intentionally, this is what makes Raman amplification work.
 - Unintentionally, it causes power transfer from one wave to another
- Tighter channel spacing actually reduces SRS effects.
 - But adding more total channels increases them.

Example max launch powers, in G.655 NZDSF fiber:


Channel Count	200GHz Spacing	100GHz Spacing	50GHz Spacing
8	15 dBm / ch	18 dBm / ch	21 dBm / ch
16	8.6 dBm / ch	11.6 dBm / ch	14.7 dBm / ch
32	2.5 dBm / ch	5.5 dBm / ch	8.5 dBm / ch
40	0.5 dBm / ch	3.6 dBm / ch	6.6 dBm / ch
80	-5.5 dBm / ch	-2.5 dBm / ch	0.5 dBm / ch

Four Wave Mixing (FWM)

- Regularly spaced signals can interact with each other, to create harmonics in other frequencies.
- The closer they're spaced, the worse the effects.
- Transmission rate independent behavior.
- Uneven channel spacing can reduce the effects.
- FWM is most prevalent in low dispersion fibers.



Four Wave Mixing Efficiency



Four Wave Mixing Examples

DWDM Channels		Fiber Chromatic Dispersion Coefficient		
Number of Channels	Channel Spacing (GHz)	2 ps/(nm·km)	5 ps/(nm·km)	10 ps/(nm·km)
		Max Signal Power (dBm)	Max Signal Power (dBm)	Max Signal Power (dBm)
8	10	-11	-6	-4
	25	-3	1	4
	50	3	7	10
	100	9	13	15
16	10	-13	-10	-6
	25	-5	-1	1
	50	0	4	8
	100	6	10	14
32	10	-14	-10	-6
	25	-6	-1	1
	50	0	4	8
	100	6	10	13

Interchannel Effects (XPM, SPM)

- Cross-Phase Modulation (XPM)
 - One wavelength of light can affect the phase of another.
 - Can cause inter-channel cross-talk on DWDM systems.
 - Also caused by mixing NRZ and Coherent systems.
 - Coherent systems actually modulate on phase, so neighboring NRZ channels cause XPM penalties in coherent channels.
 - A 100GHz (minimum) to 200GHz (best) guard band helps this.
 - High CD helps prevent XPM.
- Self-Phase Modulation (SPM)
 - Occurs when the change in signal power between a 0 and 1 is so strong that it triggers Kerr effect.
 - Low CD helps prevent SPM.

Nonlinear Effects and Effective Area

- All nonlinearities are related to the energy “density”.
 - This is why CD helps nonlinearities, it spreads the power out.
 - A “larger fiber” (technically a larger “Mode Field Diameter”) also spreads the power over a larger area, reducing peak intensity.
 - This measurement is called a fiber’s “Effective Area” (A_{eff}).
 - If not specified in the fiber specs, use MFD and $\pi * r^2$
 - Larger effective area fiber reduces all non-linear effects.
 - Some common examples:
 - Standard G.655 NZ-DSF – $50 \mu\text{m}^2$
 - LEAF or TrueWave XL NZ-DSF– $75 \mu\text{m}^2$
 - Standard G.652 “SMF28”-based NDSF – $80 \mu\text{m}^2$
 - Submarine Fiber (e.g. Corning Vascade) - $150 \mu\text{m}^2$
- One tradeoff: Larger Effective Area = Less Raman Gain


What We Transmit Over Fiber

Modulation

- At the end of the day, we live in an analog world.
 - Digital signals must still be encoded into analog waves.
 - And light is just another type of electromagnetic wave.
- The simplest form of modulation is called “IM-DD”.
 - Which stands for “Intensity Modulation with Direct Detection”.
 - The most common version is “NRZ”, or “Non-Return to Zero”.
 - Also called “Amplitude Shift Keying” (ASK).
 - Which is just a fancy way of saying “bright for a 1, dim for a 0”.
 - “Direct Detect” means only a photodiode is needed to RX.
- Historically, fiber optic systems were purely NRZ based.
 - All 10G and below optical technology is based around NRZ.

Background: Baud

- The “rate” at which you modulate a signal is the “baud”.
 - Technically defined as the “symbol rate per second”.
 - 10 Gbps means flashing bright/dim, 10 billion times/sec.
 - A.K.A. 10 GigaBaud (10GBaud)
- If you only encode 1 bit per baud, this is your bit rate.
 - Two states (bright or dim) means we represent 1 bit per symbol.
- Scaling the baud rate worked very well, to a point.
 - We built very successful networks using 10G technology.
 - And when that ran out, we built more parallel 10G links using DWDM over a single fiber.
 - But it was never enough.





So Where Do We Go From There?



- So what's a sad Internet to do when baud rates stop keeping up with our demand for cat pictures?
 - Increase the number of bits of information that can be encoded per symbol change (aka per baud)!
- There are a few methods to accomplish this.
 - Amplitude Shift Keying (ASK)
 - Have more than 2 "states", e.g. have "bright" and "really bright".
 - Phase Shift Keying (PSK)
 - Modulate on an additional property of an analog signal, by varying the "phase" of the signal over time.

Next Gen ASK: PAM2 vs PAM4


PAM2-NRZ



PAM4





PAM4



Phase Shift Keying (PSK)

- An analog signal can be represented as a sine wave.
 - A sender and receiver both agree on a specific frequency.
 - The sender transmits this frequency, but modulates the wave by “offsetting” it forwards/backwards to encode data.



Intro – Coherent Optical Technologies

- So, what exactly is “Coherent” technology?
 - The introduction of Phase Shift Keying in optical.
 - Named after it’s ability to track phase changes in optical signals (a concept called “phase coherence”).
 - Coherent provided an entirely new way to modulate signals, breaking the long-standing 10Gbps barrier.
- So how exactly does it do that?
 - By introducing a concept called a “local oscillator”.
 - AKA it uses a laser on the RECEIVE side of the signal.
 - Phase information can be computed from this reference, by comparing the received signal to the local reference laser.

Tying It All Together


- Doing Phase Shift Keying isn't as easy as it sounds.
 - With ASK all we needed was a photodiode to “see” the light.
 - But Coherent technology is built around the “DSP”.
 - “Digital Signal Processor”, an advanced purpose-built microprocessor, specifically designed for real-time processing of numeric data representing analog signals.
 - These DSPs tie all the signals together, recovering useful data.
- But Coherent technologies delivered in spades:
 - Significantly improved bandwidth (jumped from 1.6 Tbps to 9.6 Tbps)
 - Delivered true 100G optical signals, not just Nx10G signals.
 - Eliminated the need for physical Dispersion Compensation.
 - Enabled high bandwidths over massive distances.

Where Does One Go From PSK?

- Quadrature Amplitude Modulation (QAM)
 - Effectively a combination of ASK and PSK.
 - Take two amplitude modulated carriers, and send them at the same time, over the same frequency, with a phase offset (a sine and a cosine).
 - Rely on your DSP to computationally recover the signal.
- More and more complex versions can be created.
 - Adding new possible states and increasing the amount of information that can be encoded per symbol.
 - 8QAM, 16QAM, 32QAM, 64QAM, 128QAM, 256QAM....

Polarization Multiplexing

- Light is (among many other things we don't fully understand yet) a wave of electromagnetic energy, propagating through space.
- In 3-Dimensional space (e.g. a cylindrical fiber), you can send two independent orthogonal waves which propagate along a X and Y axis, which largely do not interfere with each other.
- Modern DSPs make it possible to compensate for changing fiber conditions in real time, effectively doubling bandwidth.



BPS = Polarization * Baud * Modulation

- Total transponder bandwidth is a combination of:
 - Polarization – Today dual polarization, to double capacity.
 - Baud – Higher baud needs wider channel sizes, better DACs.
 - Modulation – Higher modulation needs better OSNR levels.


Data Rate	Baud Rate	Polarities	Modulation Format	Channel Size	Raw BW (with FEC)	Efficiency (bits/s/Hz)	OSNR Required
100G	32G	2	DP-QPSK	37.5GHz	128G	2	10.5 dB
150G	32G	2	DP-8QAM	37.5GHz	192G	3	16.0 dB
200G	32G	2	DP-16QAM	37.5GHz	256G	4	19.5 dB
200G	56G	2	DP-8QAM	62.5GHz	224G	3	17.5 dB
400G	56G	2	DP-32QAM	62.5GHz	560G	5	23.0 dB
200G	64G	2	DP-QPSK	75GHz	256G	4	14.5 dB
400G	64G	2	DP-16QAM	75GHz	512G	4	21.0 dB
600G	64G	2	DP-64QAM	75GHz	768G	6	25.0 dB

Example – Acacia AC1200 DSP

	QPSK	8QAM	16QAM	32QAM	64QAM
100G	34.5 Gbaud				
150G	51.75 Gbaud	34.5 Gbaud			
200G	69 Gbaud	46 Gbaud	34.5 Gbaud		
250G		57.5 Gbaud	43.125 Gbaud	34.5 Gbaud	
300G		69 Gbaud	51.75 Gbaud	41.4 Gbaud	34.5 Gbaud
350G			60.375 Gbaud	48.3 Gbaud	40.25 Gbaud
400G			69 Gbaud	55.2 Gbaud	46 Gbaud
450G				62.1 Gbaud	51.75 Gbaud
500G				69 Gbaud	57.5 Gbaud
550G					63.25 Gbaud
600G					69 Gbaud

Baud Rates vs Higher Order Modulation

- Consider two possible ways to get to 200Gbps
 - 2 polarities * 32 Gbaud * 16QAM = 200G
 - Requires a 19.5dB OSNR, but fits in a 37.5GHz channel
 - 2 polarities * 64 Gbaud * QPSK = 200G
 - Works down to 14.5dB OSNR, but requires a 75GHz channel
- As baud rate increases, channel size requirements increase proportionally.
- If you can afford the additional spectrum, trading baud rate for modulation can significantly increase optical reach.
- But many brownfield networks with fixed channel sizes and ROADMs will not be able to fully utilize this capability.



More About Coherent

- Other Advantages of Coherent
 - Need for dispersion compensation all but eliminated.
 - Coherent DSPs eat CD for lunch - 200,000 ps/nm or more.
 - In fact, Coherent systems performance BETTER with CD.
 - Coherent can “lock on” to one specific frequency.
 - You may not need a “mux” to filter out specific channels.
 - This enables “Colorless Directionless Contentionless” ROADMs.
- But there are some major downsides too.
 - Many components, and expensive / power hungry DSP.
 - Very difficult to integrate into high-density “pluggables”.

400ZR - Coherent as a Pluggable Optic

- Historically, Coherent has always lived in “external transponders”
 - Chasing fiber “spectral efficiency”, rather than space/power efficiency.
 - Typically, in the realm of “hundreds of watts” each, while SFP/QSFP type pluggables have been designed for low-power (1-5W/ea) use only.
- But in many use cases, external transponders can be wasteful.
 - If all you’re doing with your external DCI platform is providing point-to-point bandwidth to your own routers, you can eliminate a lot of overhead by putting DWDM transponders directly onto the router (“IPoDWDM”).
- 400ZR is an attempt to bridge the gap between these two worlds.
 - On the Coherent side, a greatly stripped-down version of the functionality found in external transponders, focused on cost and power efficiency.
 - On the Pluggable side, a greatly expanded range of electrical/thermal capabilities, e.g. 12-15W for QSFP-DD or 15-20W for OSFP.
 - Plus, a well-defined standard allowing for multi-vendor interoperability.

400ZR vs OpenZR+



Reach	< 120km	> 120km
Client	400GbE Only	100-400GbE
Application	Campus, Metro, Edge DCI	DCI, Metro, Regional, Long Haul
FEC	C-FEC	<u>oFEC</u>
Max. Power Consumption	15W	~20W
Multi-vendor Interop	Yes	Yes
Standards	OIF	<u>OpenZR+ MSA</u>


The Evolution of Coherent Technologies

	Gen30		Gen60		Gen90	Gen120	
	Performance	Compact	Performance	Compact	Performance	Performance	Compact
Introduced	2010	2015	2017	2020	2020	2023 (forecast)	
Max. Bandwidth	100-200Gbps		400-600Gbps		600-800Gbps	1.2Tbps+	800Gbps
Baud Rate	~30GBaud		~60GBaud		~90GBaud	~120GBaud	
Key Technology	Phase Modulation (QPSK) Multi-level modulation (16-QAM, 8-QAM) Nyquist Shaping SD-FEC		Intelligent Modulation (PCS, Hybrid, up to 64-QAM) Real-time Telemetry Pluggable Interoperability		Nyquist Sub-carriers Electronics Innovation Co-packaging		
Si Process	65/40nm	16nm	16nm	5nm-7nm	5nm-7nm	3nm-5nm	
Examples	Infinera ICE4 Ciena WaveLogic 3e Nokia PSE-2s NEL ExaSPEED 100 Acacia AC400 Huawei OptiXtreme H3	CFP Acacia CFP2-DCO Cisco + CFP2-ACO Nokia PSE-2c Huawei OptiXtreme G6	Ciena WL Ai Nokia PSE-3s Acacia AC1200 NEL ExaSPEED TERA Huawei OptiXtreme H6	400ZR/ZR+ OpenZR+ OpenROADM 400G CFP2	Ciena WL5e Infinera ICE6 Nokia PSE-Vs Huawei OptiXtreme H7	Acacia Jannu/CIM8	

Engineering an Optical Network


Insertion Loss

- Even the best connectors and splices aren't perfect.
 - Every time you connect two fibers together, you get loss.
 - The typical budgetary figure is 0.5dB per connector.
 - Actual loss depends on your fiber connector and mating conditions.
- Insertion loss is also used to describe loss from muxes.
 - Since it is the “penalty you pay just for inserting the fiber”.
 - Some real-life examples:
 - 40-channel DWDM 100GHz Mux/Demux: 3.5dB
 - 80-channel DWDM 50GHz Mux/Demux: 9.5dB
 - Effectively just 2x 100GHz muxes (even+odd) plus an interleaver.



Balling On An (Optical) Budget

- To plan your optical network, you need a budget.
 - When an optic says “10km”, this is only a guideline.
 - Actual distances can be significantly better or worse.
 - It’s also smart to leave some margin in your designs.
 - Patch cables get bent and moved around, optic transmitters will cool with age, a fiber cut and repaired will add more loss, etc.




PC/UPC vs APC

- Beware of the different types of ferrule connectors.


- (Ultra) Physical Contact

- Blue Connectors
- PC - $< -30\text{dB}$ Back Reflection
- UPC - $< -55\text{dB}$ Back Reflection



- Angled Physical Contact

- Green Connectors
- 8° angle on the ferrule
- $< -65\text{dB}$ Back Reflection
- Incompatible with PC / UPC!
- Useful for high power applications

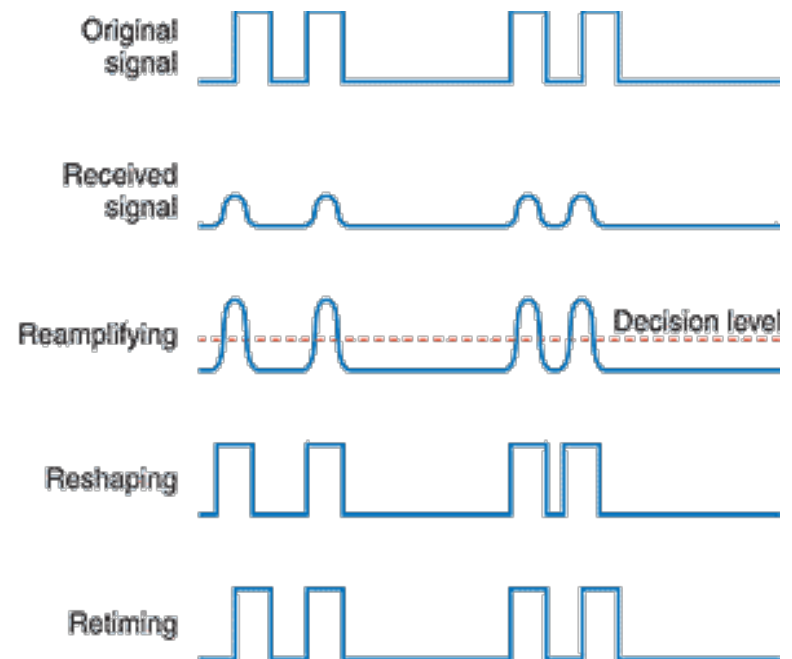


- Why? When disconnected, even UPC reflects massively.

- On a high-powered amplifier, reflections could cause damage.

Re-amplifying, Reshaping, and Retiming

- Signal Regeneration (Repeaters)
 - Different types are described by the “R’s” that they perform.
 - 1R – Re-amplifying
 - Makes the analog signal stronger (i.e. makes the light brighter)
 - Typically performed by an amplifier.
 - 2R – Reshaping
 - Restores the original pulse shape that is used to distinguish 1’s and 0’s.
 - 3R – Retiming
 - Restores the original timing between the pulses.
 - Usually involves an O-E-O conversion.




Bit Error Rates (BER)

- As optical impairments add up, links don't just "die".
 - They start taking bit errors, at progressively higher rates.
 - The probability that this will happen is the Bit Error Rate.
- For 99% confidence (100 bit error samples), test:

Date Rate	BER 10 ⁻⁹	BER 10 ⁻¹¹	BER 10 ⁻¹²	BER 10 ⁻¹³
100 Gbps	1 sec	2 min	21 min	3 hr 29 min
40 Gbps	3 sec	6 min	53 min	8 hr 47 min
10 Gbps	13 sec	21 min	3 hr 30 min	1d 10 hr 58m
1 Gbps	2 mins	3 hr 30 min	1d 10 hr 58 min	14d 13 hr 33m

OSNR(dB) and Bit Error Rates



Tools of the Trade


Optical Power Meter (or Light Meter)

- Measures the brightness of an optical signal.
- Displays the results in dBm or milliwatts (mW).
- Most light meters also include a “relative loss” function, as well as absolute power meter.
 - Designed to work with a known-power light source on the other end, to test the amount of loss over a particular fiber strand.
 - These results are displayed in dB, not dBm.
 - Frequently the source of much confusion in a datacenter, when you use the wrong mode!
 - If I had a nickel for every time someone told me they just measured a +70 signal on my fiber...




Optical Time-Domain Reflectometer (OTDR)

- An OTDR is a common tool for testing fiber.
- Injects a series of light pulses into a fiber strand.
- Analyzes the light that is reflected back.
- Used to characterize a fiber, with information like:
 - Splice points, and their locations.
 - Overall fiber attenuation.
 - Fiber breaks, and their locations (distance from the end-point).




Example OTDR Output



Feature #/Type	Location (km)	Event-Event (dB) (dB/Km)	Loss (dB)	Refl (dB)
1/N	2.3310	?? ??	0.12	
2/R	6.8035	0.91 0.203	0.64	-58.19
3/R	10.4907	0.72 0.196	1.86	-48.24
4/N	13.8639	0.70 0.206	0.06	
5/N	14.6205	0.14 0.188	-0.71	
6/N	15.9114	0.26 0.205	1.06	
7/N	17.2350	0.25 0.193	0.08	
8/E	17.4491	0.05 0.211	>3.00	>-33.55S

Overall (End-to-End) Loss: ??



Feature #/Type	Location (km)	Event-Event (dB) (dB/Km)	Loss (dB)	Refl (dB)
1/N	0.1937	0.02 0.121	-0.06 (2P)	
2/N	1.5194	0.24 0.184	-0.82	
3/N	2.8327	0.26 0.197	0.99	
4/R	6.9421	0.90 0.219	-0.21	>-46.37
5/R	10.6396	0.75 0.203	0.96	-56.69
6/R	17.2269	1.28 0.194	1.61	-61.90
7/E	17.4512	0.04 0.184	>3.00	>-34.48S

Overall (End-to-End) Loss: 5.97 dB

Question: Can I really blind myself by looking into the fiber?

Or - Beware of Big Scary Lasers



Laser Safety Guidelines

- Lasers are grouped into 4 main classes for safety:
 - Class 1 – Completely harmless during normal use.
 - Either low powered, or laser is inaccessible while in operation.
 - Class 1M – Harmless if you don't look at it in a microscope.
 - Class 2 – Only harmful if you intentionally stare into them
 - Ordinary laser pointers, supermarket scanners, etc. Anyone who doesn't WANT to be blinded should be protected by blink reflex.
 - Class 3 – Should not be viewed directly
 - Class 3R (new system) or IIIA (old system)
 - Between 1-5mW, “high power” Internet purchased laser pointers, etc.
 - Class 3B (new system) or IIIB (old system)
 - Limited to 500mW, requires a key and safety interlock system.
 - Class 4 – Burns, melts, destroys Alderaan, etc.

Laser Safety And The Eye

- Networking lasers operate in the infrared spectrum
 - Infrared can be further classified as follows:
 - IR-A (700nm – 1400nm) – AKA Near Infrared
 - IR-B (1400nm – 3000nm) – AKA Short-wave Infrared
 - Laser safety levels are based on what can enter the eye.
 - Remember, the human eye didn't evolve to see infrared.
 - The cornea actually does a good job of filtering out IR-B light.
 - So IR-B has much higher safety limits than visible light.

Max power (continuous, without auto-shutdown features) for IR-B:

Class 1	Class 3R	Class 3B	Class 4
< 10 dBm	< 17 dBm	< 27 dBm	> 27dBm

Optical Networking and Safety


- Routers
 - Essentially every single-channel laser that can be connected to a router is a Class 1 or Class 1M laser.
 - Even “long reach” 200km+ optics are no exception:
 - A multi-lane optic can have the highest output, e.g. 40G LR4 = 8mW
- Optical Amplifiers
 - Can easily have output powers of 3R (metro) or 3B (long-haul).
 - Raman amplifiers are almost always Class 4.
 - But they all have Automatic Power Reduction/Shutdown too.
- DWDM Equipment
 - Total output power is the sum of all muxed input signals.
 - This can put the total output power into the 3B territory even without amplification, and often has no auto-shutdown feature.

Optical Networking and Safety

- So should I be wearing safety goggles in the colo?
 - Generally speaking, your standard client optics are always Class 1 (completely safe under all conditions).
 - Even on amplified/DWDM systems, light rapidly disperses as soon as it leaves the fiber and travels through air.
 - Wavelengths above 1400nm are IR-B, and are mostly blocked by the human eye. Most high power optics and long-reach systems are in this range.
 - High-power systems are legally required to have auto-shutdown safety mechanisms if they detect a cut.
- But, don't hold a DWDM mux directly to your eye.
- And be extra careful with a fiber microscope.

Why Look Into The Fiber Anyways?

- Can you even see the light at all?
 - No, the human eye can only see between 390 – 750nm.
 - No telecom fiber signal is directly visible to the human eye.
- But, I looked at 850nm and I saw red?
 - What you're seeing are the sidebands of an imperfect signal generation, not the main 850nm signal itself.
 - Many digital cameras can see infrared.
 - One trick to check for light in a fiber is to hold it up to your camera phone.
 - You can try this on your TV's remote control.
 - Except newer/nicer ones filter IR, for picture quality. iPhone started blocking IR as of 4S/5.



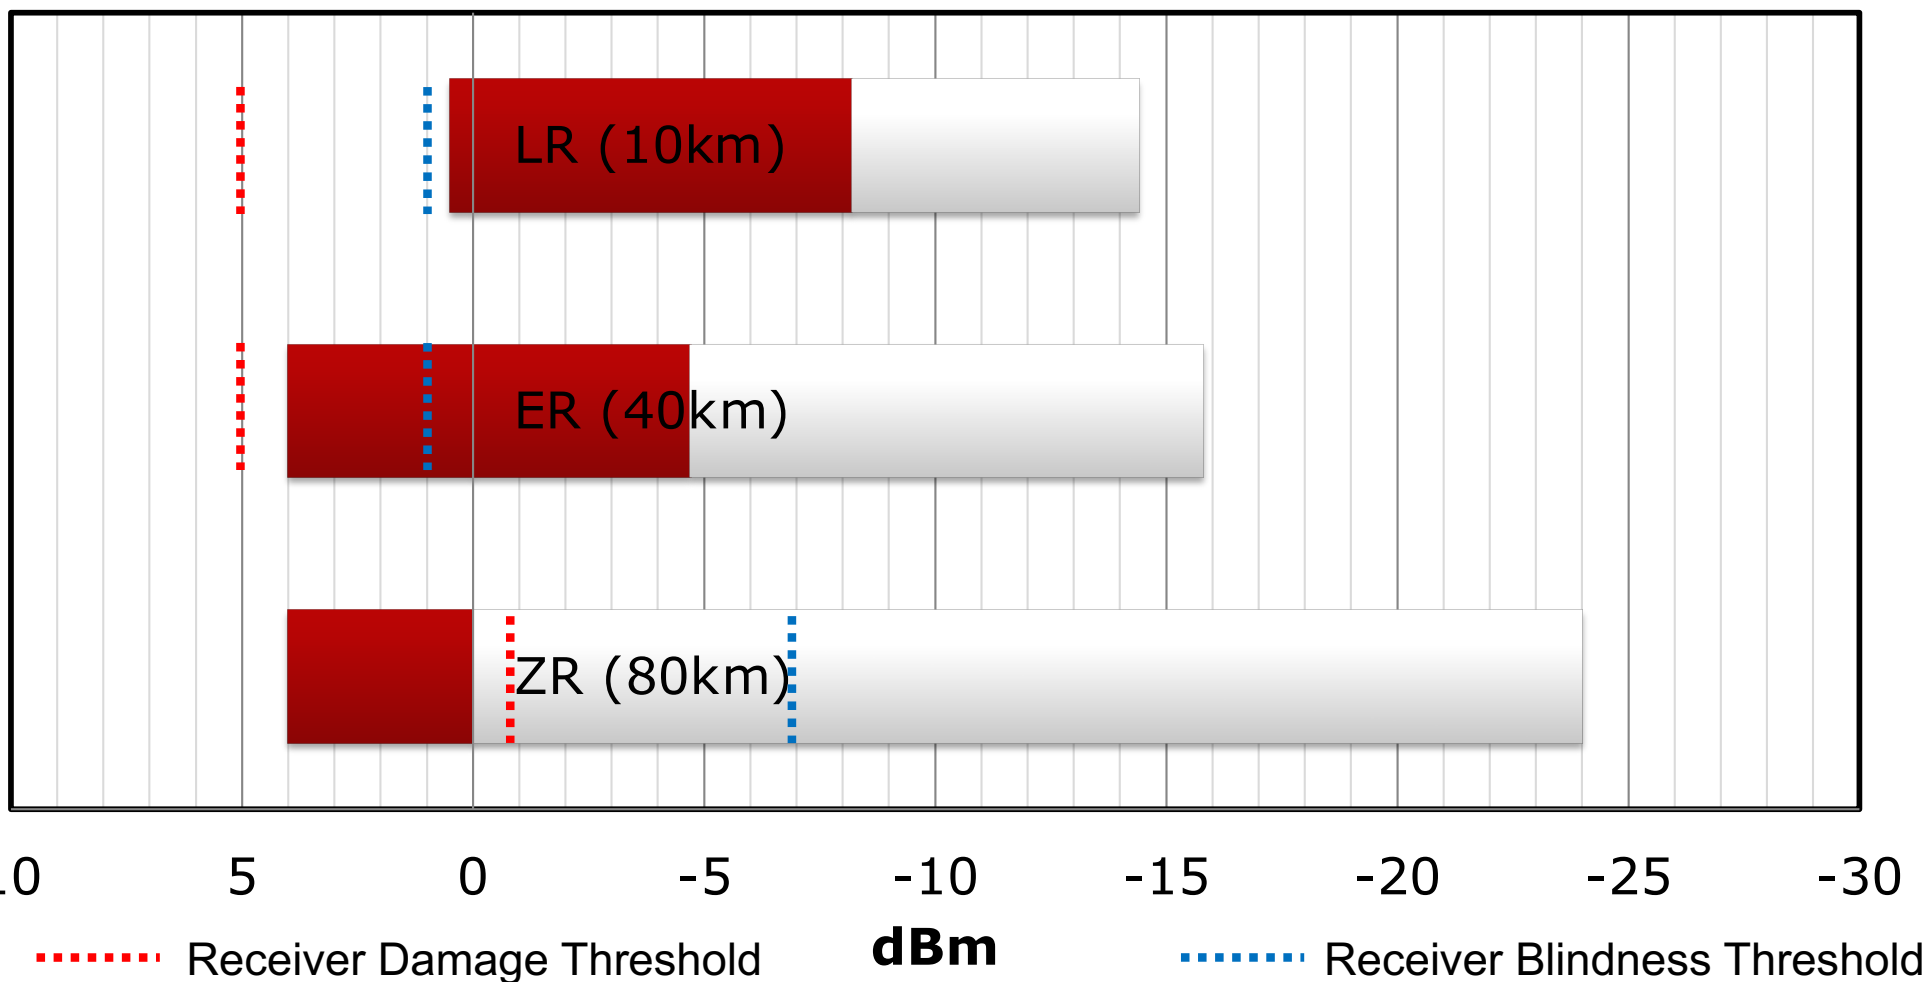
Question: Can optical transceivers be damaged by over-powered transmitters?

Damage by Overpowered Transmitters?

- Actually, most optics transmit at roughly the same power.
 - The typical output of 10km vs 80km optics are within 3dB.
- Long reach optics achieve their distances by having more sensitive receivers, not stronger transmitters.
 - 80km optics may have a 10dB+ more sensitive RX than 10km.
 - These sensitive receivers are what are in danger of burning out.
- There are two thresholds you need to be concerned with.
 - Saturation point (the receiver is “blinded”, and takes errors).
 - Damage point (the receiver is actually damaged).
 - The actual values depend on the specific optic.
 - But generally speaking, only 80km optics are at risk.

Tx and Rx Optical Power Ranges

■ Tx Window ■ Rx Window





Question: Do I really need to be concerned about bend radius?

Is Bend Radius Really A Concern?

- Yes, bend radius is a real issue.
 - Remember that total internal reflection requires the light to hit the cladding below a “critical angle”.
 - Bending the fiber beyond its specified bend radius causes the light to “leak” out.
 - There are “bend insensitive” fibers, though they usually trade some level of performance for this.
 - These are pretty useful in datacenter applications, when humans don’t do the right thing.



Practical Bend Radius Examples (SMF)



Question: Can two transceivers on different wavelengths talk to each other?

Can You Mismatch Transceiver Freqs?

- Between certain types of optics, yes.
 - All optical receivers have wideband photodetectors.
 - Laser receivers “see” everything between 1260nm – 1620nm.
 - But they won’t be able to see a 850nm LED, for example.
 - Coherent receivers can even “lock on” to one specific frequency.
 - Many DWDM networks are built around this premise.
 - By using one wavelength going A->B and other going B->A, you can achieve a bidirectional system over a single fiber strand.
 - The DWDM filters (muxes and OADM)s provide hard cut-offs of certain frequencies, but the transceivers can receive any color.
 - The only “gotcha” is optical power meters will be wrong.
 - A meter that is calibrated to read a 1310nm signal will see a 1550nm signal just fine, but its power reading will be a few dB off.

Can You Mismatch Transceiver Freqs?


- You can also mismatch frequencies for added reach.
 - You can achieve nearly as much distance with an LR/ER pair (1310nm 10km / 1550nm 40km) as with an ER/ER pair.
 - The ER transmits at 1550nm, which has a lower rate of attenuation.
 - Around 0.2dB/km vs 0.35dB/km, depending on fiber type.
 - So the LR side receives a much stronger signal than the ER side.
 - And the ER optic has a much greater RX sensitivity than the LR.
 - So it will be able to hear the 1310nm signal much better than an LR optic would in the same position.- Result:
 - You may only *need* a long reach optic on one side.

Question:

Do I Really Need to Clean the Fiber?

Do I Really Need to Clean the Fiber?

- Dirt can actually DAMAGE the connector permanently.
 - A mating force of 2.2lb, over a 200 μ m surface area...
 - Results in 45,000 lbs per square inch of pressure.
 - This can permanently pit and chip your fiber cables!
- Buy a cheap cleaning kit!



Other Misc Fiber Information

How Fast Does Light Travel In Fiber?

- Ever wondered how fast light travels in fiber?
 - The speed of light is 299,792,458 m/sec
 - SMF28 core has a refractive index of 1.4679
 - Speed of light / 1.4679 = 204,232,207 m/sec
 - Or roughly 204.2 km/ms, or 126.89 miles/ms
 - Cut that in half to account for round-trip times.
 - So, approximately 1ms per 100km (or 62.5 miles) of RTT.
- Why do you see a much higher value in real life?
 - Remember, fiber is rarely laid in a straight line.
 - It is often laid in rings which take significant detours.
 - Dispersion compensation can add extra distance too.

Send questions, comments, complaints to:

Richard A Steenbergen <ras@petabitscale.com>